TIMING BINARY AND MILLISECOND PULSARS

J. H. Taylor
Joseph Henry Laboratories and Physics Department,
Princeton University,
Princeton, NJ 08544 USA

ABSTRACT. This tutorial lecture outlines the essentials of making and analyzing timing observations of binary and millisecond pulsars. Many of the necessary procedures are, of course, similar to those used for slow, single pulsars. However, achieving the best attainable timing accuracy for millisecond pulsars requires much tighter tolerances on observing and data-processing procedures, and the presence of orbital motion qualitatively complicates both data acquisition and analysis. In return for adequate diligence in these matters, an observer gains access to a much richer set of measurables than are available for ordinary pulsars. Binary pulsars provide the only experimental handles on the masses of non-accreting neutron stars, and the parameters of their orbits and characteristics of their companion stars provide unique clues on the origin and evolution of the systems. In the most favorable circumstances, significant tests of fundamental physical laws are even possible. Timing data on millisecond pulsars turns out to be valuable for other reasons as well: because they are extremely stable clocks, these objects provide exquisite tools for a wide range of studies in fundamental astrometry, cosmology, gravitation physics, and metrology.

Don Backer has already given us a sound primer on radio pulsar timing techniques (Backer 1989). To review briefly, I will remind you that a typical observing procedure uses a setup like the one diagrammed in Figure 1, used by my colleagues and me at Arecibo Observatory over the past several years (Rawley, Taylor, Davis, & Allan 1987; Rawley, Taylor, & Davis 1988). Incoming signals from the telescope are amplified, converted to intermediate frequency, and passed through a "filter bank" spectrometer which analyzes the total accepted bandwidth into channels narrow enough that the observed pulse widths are not dominated by dispersive smearing. Synchronous signal averaging is used to accumulate estimates of a pulsar's average waveform in each of the spectral channels, using electronics under control of a small computer and accurately synchronized with the Observatory's time and frequency

H. Ögelman and E. P. J. van den Heuvel (eds.), Timing Neutron Stars, 17–26.

standard. A programmable synthesizer, whose output frequency is updated once a second in a phase continuous manner, compensates for variable Doppler shifts due to motions of the pulsar and the observatory. Integrated pulsar waveforms are recorded once every few minutes, together with appropriate time tags.

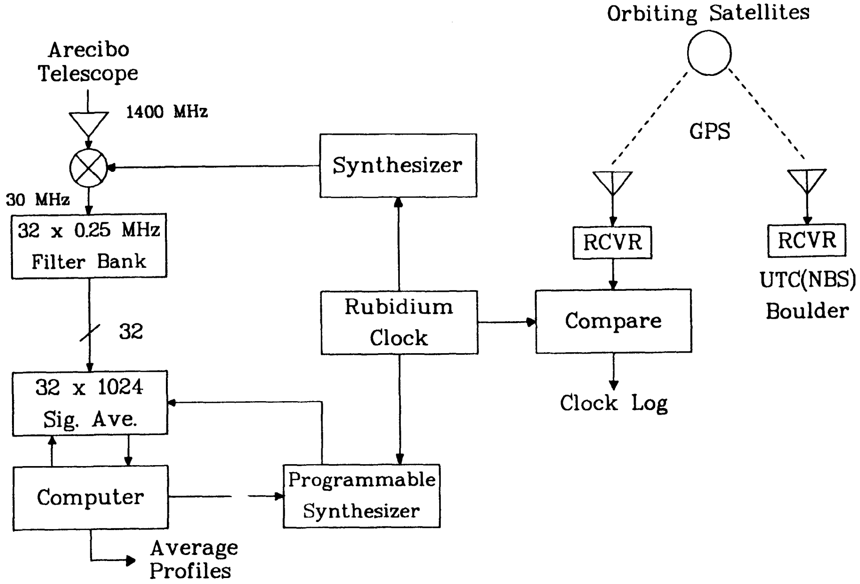

Figure 1: Block diagram of a pulsar timing system used for binary and millisecond pulsars at Arecibo Observatory.

Analysis of recorded profiles usually follows a series of steps similar to those listed in cookbook Recipe 1 on the next page. The first step after observing involves "template matching" to determine the phase of each recorded profile relative to the start of its integration. This process can be carried out by cross-correlation in the time domain, or by an equivalent procedure in the Fourier transform domain. There may be reasons of operational convenience to prefer one of these methods over the other; more importantly, there are significant reasons to prefer the Fourier transform approach when the signal-to-noise ratio is high and the achievable timing uncertainty is much smaller than the interval at which the waveform has been sampled (Rawley 1986). (These conditions are true, for example, in our observations of the millisecond pulsar 1937+21 at Arecibo.) After a phase offset has been determined, the corresponding time delay is added to the start time of the integration to yield a topocentric time of arrival, or TOA.

Recipe 1. Pulsar timing observations and analysis.

1. Observe average profiles at known UTCs.
2. Determine topocentric TOA's by template matching.
3. If necessary: remove dispersive delays, sum channels.
4. Apply clock corrections to yield TDT.
5. Transform to TDB at solar system barycenter.
6. Bootstrap a timing model (see Recipe 2 or 3).
7. Repeat from step (1) as often as possible.

$\Downarrow$

8. Output: α, δ, ϕ_0, P, $\dot{P}$; possibly also μ_α, μ_δ, $\ddot{P}$, and orbital parameters.

For a given pulsar, the optimum template or "standard profile" for use in the matching procedure is a high signal-to-noise version of the profile, generally observed with the same equipment but over a longer integration time. For pulsars too weak to give adequate signal-to-noise ratio in a single filter bank channel, the data in all channels can be summed (after shifting phases to account for the different dispersion delays at each frequency) and a single equivalent TOA determined.

An accurate time transfer system, using signals from satellites in the Global Positioning System, allows us to express all TOA's in terms of Coordinated Universal Time as maintained and distributed by the US National Bureau of Standards, or in terms of any of a number of other high-precision atomic time standards kept at national timekeeping laboratories around the world. An appropriately weighted average of these standards represents the best currently available approximation to an ideal implementation of terrestrial dynamical time, or TDT, and is the reference standard of choice for the most exacting observations.

A model capable of predicting pulse arrival times is most conveniently formulated in an inertial reference frame, for which the solar system barycenter is an adequate approximation. The necessary general relativistic transformation from TDT to barycentric dynamical time, or TDB, must be done with care (see Backer & Hellings 1986). The transformation depends on accurate knowledge of the pulsar's celestial coordinates, α and δ, and unambiguous numbering of the received pulses requires good knowledge of the pulsar period P and spin-down rate, $\dot{P}$. Since these quantities may be poorly known at the beginning—they are, after all, among the parameters one is trying to measure—the process necessarily requires an iterative "bootstrap" approach like that outlined in Recipe 2.

Recipe 2. Bootstrap procedure for timing single pulsars.

1. Collect data at various intervals—say 2 m, 10 m, 2 h, 1 d, 10 d, 3 mo, ...
2. Obtain least-squares estimates of P, α, δ, $\dot{P}$ (or a subset).
3. Inspect residuals for non-random appearance ($\Rightarrow$ evidence of mis-numbering of pulses across gaps in data).
4. Repeat, possibly adding or subtracting integers in the pulse numbering scheme, until *you've got it!*

One of the principal reasons that pulsar timing data can yield so much information is that the rotation phase of a non-accreting neutron star can be measured to within a small fraction of a turn, and predicted with similar accuracies over intervals as long as years. Therefore phase-coherent solutions are possible, spanning many years of observations and perhaps 10^8–10^{11} rotations of the neutron star. In such solutions the integer pulse number corresponding to each observed TOA is determined unambiguously.

When timing observations are begun for a new pulsar, the bootstrap procedure to reach the desirable phase-connected state of affairs is facilitated if data are collected at various intervals in rough geometric progression, with increments no more than a factor of 10 (see Recipe 2). The largest gaps between available TOA's will then be small enough that one should be able to extrapolate across them without pulse numbering ambiguities, or at worst with only a small range of integer values permitted. A least-squares solution can then be carried out to determine improved values of P and, when enough data are available, also α, δ, and $\dot{P}$. A plot of the post-fit residuals from the solution gives a clear indication of whether the pulses have been numbered correctly. Some trial and error may be required, but the correct solution is easily recognized when attained.

Three examples of post-fit residual plots are illustrated in Figure 2. These observations, made with the NRAO 92 m telescope at Green Bank by Dewey *et al.* (1988), were concentrated in nine observing sessions between January 1985 and May 1987. Each observing session lasted several days, and each of about 75 pulsars was observed for 5 to 10 two-minute integrations on one or more days in a session. Thus the data set for a given pulsar generally consists of TOA's with spacings of 2–20 minutes, 1–4 days, 15 days, and several months to 2.3 years. The residuals in the examples amount to a few milliperiods or less, and appear to be essentially random. Solutions like these typically determine P to 11 or 12 significant figures, $\dot{P}$ to an uncertainty $\sim 10^{-18}$, and the pulsar coordinates to within about $0.1''$.

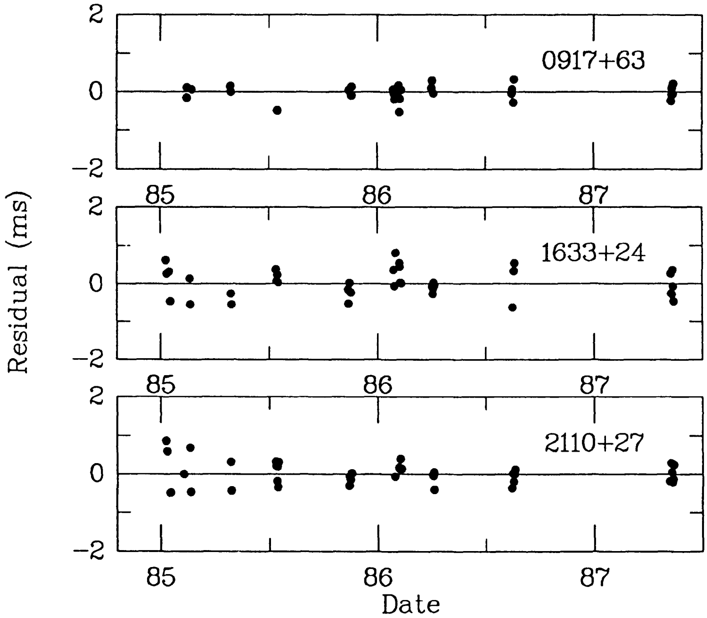

Figure 2: Plots of the post-fit residuals for timing observations of three pulsars (after Dewey *et al.* 1988). For clarity only a single datum is plotted for each observing day.

For binary pulsars the procedure is similar in concept, but a number of additional parameters are involved, including five elements of a Keplerian orbit and possibly observable general relativistic effects as well. In formulating the timing model an additional relativistic transformation is required, dependent on the orbital parameters, to convert from TDB to proper time at the pulsar. Since the orbital parameters are not known (even approximately) at the outset, the bootstrap procedure to attain a phase connected solution is usually much more difficult than with single pulsars.

A workable scheme that my colleagues and I have used to obtain phase connected solutions for 8 binary pulsars is summarized in Recipe 3. When initial timing observations suggest that the period of a pulsar is variable, we proceed to collect as many TOA's as possible. The data are separated into blocks no longer than $\sim$ 20 m and periods are computed for each block, reduced to the solar system barycenter. In principle a plot of these periods as a function of time should trace out the orbiting pulsar's velocity curve. In practice, however, the data are likely to be coarsely and irregularly sampled, and recognizing the orbital period may be very difficult.

An algorithm for determining the orbital period in such circumstances is out-

Recipe 3. Bootstrap procedure for timing binary pulsars.

1. Collect data: get TOA's and pulsar periods P_i at many epochs.
2. Determine approximate orbital period, P_b (see Recipe 4).
3. Incoherent solution: use P_i's to obtain least-squares estimates of P, $a_1 \sin i$, e, T_0, P_b, ω.
4. Phase-connected solution: use TOA's to obtain least-squares estimates of P, $a_1 \sin i$, e, T_0, P_b, ω, α, δ, $\dot{P}$, and possibly relativistic parameters.

lined in Recipe 4. Its equations contain nothing about Kepler's laws or the expected shapes of orbital velocity curves; the procedure works simply by seeking the smoothest possible dependence of observed period on computed orbital phase, consistent with the observations. When an approximate orbital period has been found, we use the (P_i, t_i) pairs to carry out a least-squares solution for the orbital elements (step 3, Recipe 3). Since there is no attempt to "connect phase" between the various independent measurements, pulse numbering ambiguities are not an issue. The orbital elements determined from this solution are then used as input values for a phase connected solution, which yields optimized estimates for all of the interesting parameters.

Recipe 4. Algorithm for finding a binary pulsar's orbital period, P_b.

1. Obtain solar-system barycentric periods P_i at many epochs t_i.
2. Set trial P_b to minimum reasonable value.
3. Compute orbital phases, $\phi_i = \mathrm{mod}(t_i/P_b, 1.0)$.
4. Sort list of (P_i, t_i, ϕ_i) triplets in order of increasing ϕ.
5. Compute $s^2 = \sum(P_j - P_{j-1})^2$, omitting any terms for which $\phi_j - \phi_{j-1} > 0.1$. ($j$ is the index after sorting).
6. Increment $P_b := [1/P_b - 0.1/(t_{max} - t_{min})]^{-1}$.
7. Repeat from (3) until maximum reasonable P_b is reached.
8. Choose the P_b that yielded the smallest normalized s^2.

The most complete and well documented model for analyzing binary pulsar timing data is that of Damour & Deruelle (1986), based on a highly accurate solution they developed for the general relativistic two-body problem. In addition to the usual parameters for single pulsars and five Keplerian orbital parameters, my implementation of their model includes as many as six "post-Keplerian" parameters. As an example of its use, Table 1 presents the latest list of measured orbital parameters of the PSR 1913+16 system. Further details may be found in Weisberg & Taylor (1984) and references therein.

Table 1. Orbital parameters of PSR 1913+16.

Keplerian orbital elements	
Projected semimajor axis	$a_p \sin i = 2.341769(24)$ light sec
Eccentricity	$e = 0.6171311(14)$
Orbital Period	$P_b = 27906.981644(12)$ s
Longitude of periastron	$\omega_0 = 178.86372(29)$ deg
Julian ephemeris date of periastron, and reference time for P_b and ω_0	$T_0 = 2442321.4332075(9)$
Post-Keplerian parameters	
Mean rate of periastron advance	$\langle\dot{\omega}_0\rangle = 4.22660(10)$ deg y^{-1}
Gravitational redshift and time dilation	$\gamma = 4.302(24)$ ms
Orbital period derivative	$\dot{P}_b = (-2.40 \pm 0.04) \times 10^{-12}$ s s^{-1}

Timing observations of binary pulsars have provided information that has been extremely useful in delimiting the range of possible evolutionary schemes for pulsars and neutron stars. I'll have no time in this lecture to discuss these matters further; some of the important ideas have already been outlined for us by Ed van den Heuvel (1989), and I have no doubt we'll be hearing more about them over the next few days. For a more detailed summary of results and some very recent work, see Taylor (1987) and Taylor & Dewey (1988).

In the remainder of this talk I will deal with some special issues relevant to timing millisecond pulsars. For these objects it is possible to measure TOA's with observational uncertainties well under a microsecond; consequently, to avoid contaminating good data by sloppy handling, and to extract the maximum possible information content, one must take extraordinary care in the data analysis. As an example of the quality of data obtainable, and to illustrate one particular type of observational problem that must be dealt with, Figure 3 illustrates post-fit residuals for the data on PSR 1937+21 that we have been acquiring at Arecibo Observatory since October 1984.

Both portions of Figure 3 are based on exactly the same data, analyzed in

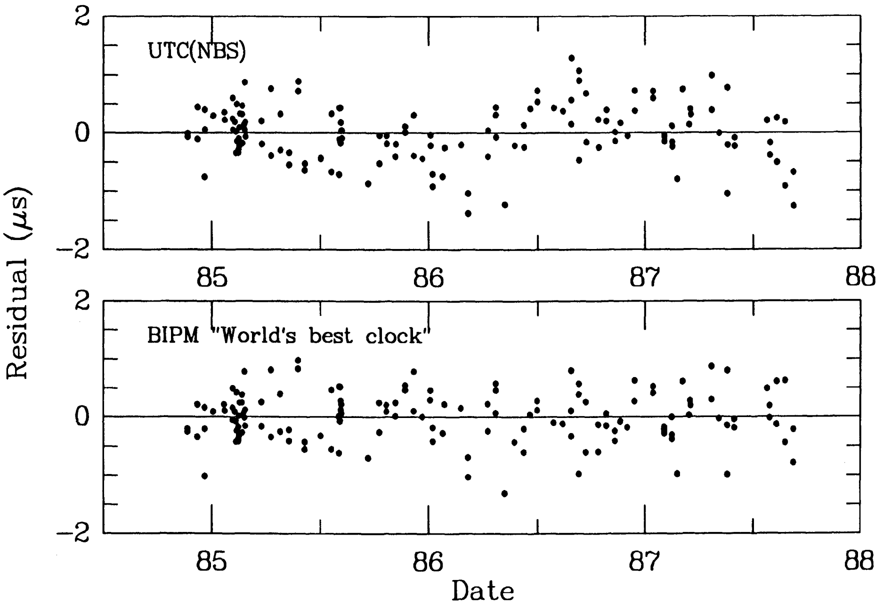

Figure 3: Post-fit arrival time residuals for PSR 1937+21 relative to UTC(NBS) (top) and to the BIPM "World's best clock" (bottom).

exactly the same way, except for using different time standards. For the upper portion of the Figure we used UTC(NBS), the version of Coordinated Universal Time distributed by the US National Bureau of Standards. This time scale, believed to be stable to within a few parts in 10^{14} over timescales of a year or so, is apparently not quite stable enough to serve as an adequate reference for this experiment. Neither, incidentally, is *any other* of the independent time standards maintained at timekeeping laboratories around the world. Some are a bit worse, and one or two appear to us to be a bit better—especially the one maintained by the Physikalisch-Technische Bundesanstalt in the Federal Republic of Germany—at least within our limited experience since late 1984. Significantly, we think, most of the slow, systematic wandering of the PSR 1937+21 residuals in the top panel of Figure 3 disappears when we use a weighted mean of of the world's best clocks (computed by the Bureau International de Poids et Measures in Paris) as the reference standard. Residuals from a such fit are shown in the lower panel of Figure 3.

People in the time-and-frequency business assure me that they have not been not standing still, and that better reference standards will be available Real Soon Now. In the meantime, we astronomers should not overlook the possibility of creing our own time standard by comparing TOA's from a number of millisecond pulsars. At present the next best pulsar with a substantial quantity of accumulated timing data is PSR 1855+09, which shows random residuals of about 2 μs rms after 2.3 years. Unfortunately, timing measurements with uncertainties in the sub-microsecond range have not yet been achieved for any pulsar besides PSR 1937+21—except for the very recently discovered PSR 1957+20, which Andy Fruchter will speak about shortly.

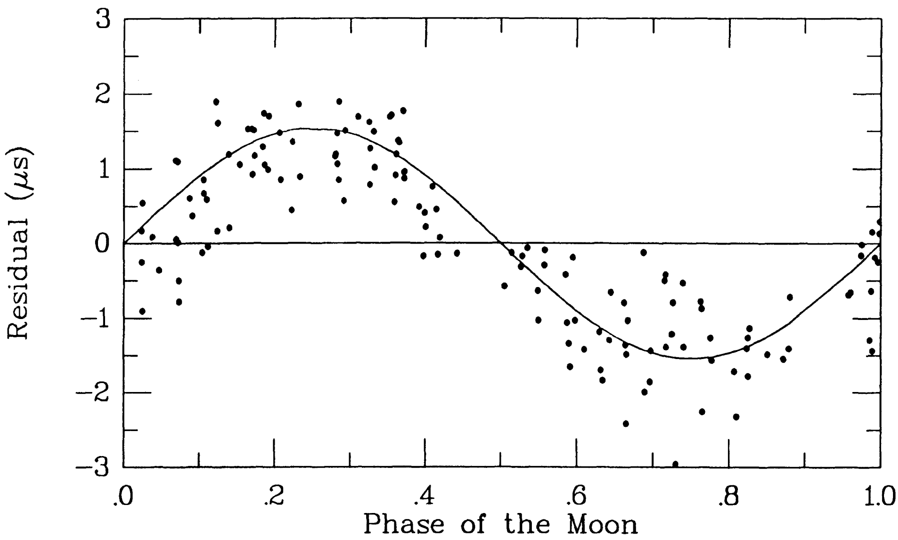

Figure 4: Post-fit arrival time residuals for PSR 1937+21, for a solution in which the monthly changes in gravitational redshift of terrestrial clocks was intentionally ignored. The smooth curve illustrates the omitted term in the time transformation.

I have already mentioned that in order to model the expected TOA's of any pulsar with the accuracies implicit in Figure 3, extraordinary care is required in the analysis procedure. An example of what can go wrong when a small effect is ignored is presented in Figure 4. For this solution we intentionally omitted the monthly cycle in the relativistic transformation from TDT to TDB. (All clocks on Earth run slightly slower when the Moon is full, because the Earth is deeper in the solar gravitational potential; furthermore, the magnitude of the Earth's total velocity changes with lunar phase, causing changes in time dilation.) In Figure 4 the residuals resulting from this omission are plotted as a function of phase of the moon, together with a curve corresponding to the omitted term. Interestingly, one can use this effect to perform Einstein's gravitational redshift test of the equivalence principle. For the ratio of observed to expected monthly redshift effect, we obtain 1.04 ± 0.06 (Gelfand 1988).

There is little doubt that further examples of binary and millisecond pulsars will be found, and that timing observations of them—as well as continued timing observations of those already known—will continue to pay very worthwhile dividends.

REFERENCES

Backer, D. C. 1989, this volume.

Backer, D. C., & Hellings, R. W. 1986, *Ann. Rev. Astron. Astrophys.*, **24**, 537.

Damour, T., & Deruelle, N. 1986, *Ann. Inst. H. Poincaré (Physique Théorique)*, **44**, 263.

Dewey, R. J., Taylor, J. H., Maguire, C. M., & Stokes, G. H. 1988, *Astrophys. J.*, in press.

Gelfand, B. Y. 1988, unpublished Princeton University "Junior Paper."

Rawley, L. A. 1986, Princeton University Ph.D. thesis.

Rawley, L. A., Taylor, J. H., & Davis, M. M. 1988, *Astrophys. J.*, **326**, 947.

Rawley, L. A., Taylor, J. H., Davis, M. M., & Allen 1987, *Science*, **238**, 761.

Taylor, J. H., 1987, *The Origin and Evolution of Neutron Stars*, ed. D. J. Helfand and J.-H. Huang, Dordrecht: Reidel, p. 383.

Taylor, J. H., & Dewey, R. J., *Astrophys. J.*, in press.

van den Heuvel, E. P. J. 1989, this volume.

Weisberg, J. M., & Taylor, J. H. 1984, *Phys. Rev. Letters*, **52**, 1348.