PULSE TIMING AND NEUTRON STAR STRUCTURE

Paul E. Boynton
University of Washington

The implementation of pulse timing analysis as a probe of neutron star structure is reviewed for both pulsars and X-ray pulsators. Current results are particularly significant for the Crab pulsar.

I. INTRODUCTION

Pulsar timing noise is generally interpreted as a stochastic variation in the rotation rate of the neutron star and its associated magnetosphere, and as such must arise from some fluctuating torque. Broadly speaking, there are only two possible sources of torque: external torque (radiation torque for pulsars), and internal torque possibly stemming from the nature of the coupling between the neutron star crust and a faster-spinning superfluid component in the interior. Whether restless behavior of the spin rate is due to either external or internal torques, the response of the system (to the extent that nonrigid-body behavior is revealed) provides a window to the internal structure of neutron stars, or equivalently, can lead to constraints on the equation of state of neutron star matter, and neutron star mass (see Lamb 1979).

In this connection I would emphasize the importance of the X-ray pulsators, which provide a similar diagnostic possibility. For accreting neutron stars secular spin-up suggests a source of internal torques similar to that present in pulsars, but in this case fluctuating external torques may result from the process of matter accretion. Detailed studies of several pulsating X-ray sources reveal pulse periods which are not smoothly decreasing, and in fact are much noisier than that of the Crab. Each source (including Hercules X-1) shows period fluctuations even to the extent of occasional episodes of spin-down. Again, interpreted directly as variations in the angular velocity of the neutron star crust, such fluctuations may provide considerable insight to the nature of accretion torques as well as the constitution of neutron star interiors as detailed by Lamb, Pines and Shaham (1978).

In this paper we review the application of such ideas to the Crab

W. Sieber and R. Wielebinski (eds.), Pulsars, 279–290.

https://doi.org/10.1017/S0074180900093037 Published online by Cambridge University Press

pulsar and to Hercules X-1, and note that rigid-body behavior is indicated over a large bandwidth. In the case of the Crab this discovery implies severe difficulties for the well-worn two-component neutron star model of Baym, Pines, Pethick and Ruderman (1969).

II. PROBING NEUTRON STAR STRUCTURE

To understand how internal structure may be inferred from observations of pulsar timing noise, consider a rotating system having an effective moment of inertia whose value depends on the timescale, t, of the applied impulse. For instance, if a putative neutron star crust and superfluid are coupled over some characteristic time, τ, then for torques applied to the crust-magnetosphere component, $I_{eff}(t/\tau) = I_c$ for $t \ll \tau$ and $I_{eff}(t/\tau) = I_c + I_n = I_s$ for $t \gg \tau$ (where I_n refers to the superfluid neutron interior, and the subscript c denotes crust). The dynamical response of the star mirrors this change in I_{eff}. Since $N = I_{eff}\,\dot{\nu}$ then $\Delta\dot{\nu}_t \simeq \Delta N_t / I_{eff}(t/\tau)$, where $\Delta\dot{\nu}_t$ is an increment in $\dot{\nu}$ on timescale t. Thus a stochastic torque perturbation (normalized to unit variance) will produce a variance in $\dot{\nu}$ which depends rather strongly on the impulse timscale: $\langle(\Delta\dot{\nu})_t^2\rangle \simeq I_{eff}^{-2}(t/\tau)$.

A more useful description is provided by a statistical treatment couched in terms of power density spectra. For simplicity consider broad-band torque noise whose power spectrum, $P_N(\omega)$, is flat (white torque noise). The response of the star is characterized by the power spectrum of fluctuations in $\dot{\nu}$, $P_{\dot{\nu}}(\omega)$. The common expectation for this response follows from the viscously coupled two-component model of Baym et al. (1969), for which they write: $2\pi I_c\dot{\nu} = -N + 2\pi I_c(\nu_n - \nu)/\tau$, and $2\pi I_n\dot{\nu}_n = -2\pi I_c\,(\nu_n - \nu)/\tau$. The power spectrum, $P_{\dot{\nu}}(\omega)$, is directly computed as the square of the Fourier transform of $\dot{\nu}(t)$ arising from a torque perturbation N(t). Thus from these two coupled equations,

$$P_{\dot{\nu}}(\omega) = [((1 - Q)^2\,\omega_d^2 + \omega^2)/(\omega_d^2 + \omega^2)]\;P_N(\omega) \qquad (1)$$

where $(1-Q)^2 = (I_c/I_s)^2$, $\omega_d = 1/\tau_d$, and $\tau_d = Q\tau$ is the exponential decay time following a δ-function torque impulse.

Thus for white torque noise, $P_N(\omega)$ = constant, the response as measured by the angular acceleration spectrum has the distinctive form shown in figure 1. From equation 1 the ratio of low frequency to high frequency power densities depends specifically on neutron star structure, $P(low)/P(high) = (1-Q)^2 = (I_c/I_s)^2$. Since the ratio of crust to star moments follows from the equation of state of neutron star matter and the star mass, P(low)/P(high) is the kind of diagnostic probe which we seek.

To carry out this program, a fundamental requirement is to be able to examine the power spectra of both stimulus (torque) and response (angular acceleration). Only modest (one octave) resolution is required to observe the behavior in figure 1, but the widest possible range of frequency, ω, must be examined in order to separately determine $P_N(\omega)$; and even that task is ambiguous unless it has a smooth, monotonic form. Fortunately, our work indicates that $P_N(\omega)$ is a simple power-law for the sources we have investigated.

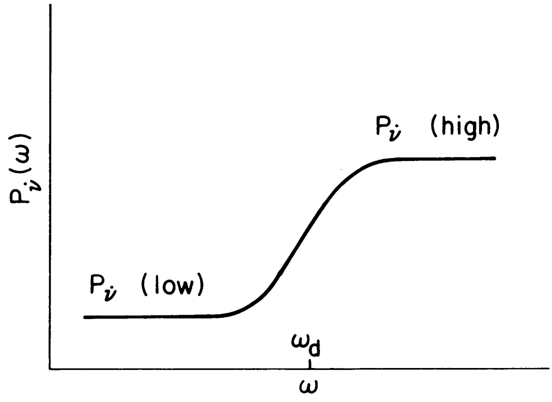

Figure 1: Power density "step" characteristic of viscously coupled crust-core model driven by white torque noise

At this point the program seems fairly clear. We calculate power spectra of $\dot\nu$ fluctuations from pulse timing observations, infer the form of $P_N(\omega)$, then perform a least squares fit of the "structure function" given by eq. 1 to the observed spectrum $P_{\dot\nu}(\omega)$. In doing so, depending on which parameters are involved in the relaxation, we either provide minimum variance estimates of the relevant structure function parameters Q and τ_d, or pose a hypothesis test in which a particular pair of Q, τ_d values is found to be rejected at some level of significance. However, rather than proceed directly to these goals, we pause to review several of the difficulties which must be confronted in recognizing data suitable for such an analysis, and also to understand the problems implicit in the calculation of the type of power spectra we have encountered.

III. SKINNING THE ONION

Arriving at a power spectrum of timing fluctuations from raw timing data is a complex task whose major elements are outlined in the following paragraphs. To begin, a provisional ephemeris is chosen with sufficient care to enable pulses to be "folded" over intervals containing many pulses without smearing. The resulting "average" pulses are superposed by maximizing their cross-correlation, thereby forming a master pulse. By cross-correlating the master pulse with successive average pulses, a sequence of pulse arrival times is established, which then provides a means for improving the ephemeris and master pulse through iteration.

Carrying out the stages of refinement of the provisional ephemeris is something like removing the layers of an onion. The initial choice may be

https://doi.org/10.1017/S0074180900093037 Published online by Cambridge University Press

completely phenomenological, since the only requirement is that the pulse should not be smeared appreciably when folded over some "local" interval. Subsequent corrections are modeled around several successive coordinate transformations, first to transform the observer to a local inertial frame -- the solar system barycenter. Then for a pulsar or X-ray pulsator in a binary system, there is a second transformation to the source frame through a source orbital ephemeris. Having peeled away two "layers" of relative motion between source and observer and thereby arriving at pulse emission times, one generally finds a further transformation is required to "remove" the secular spin-up (X-ray pulsator) or spin-down (pulsar) of the source.

Now, can the remaining variation in emission times be treated as noise? In the case of the Crab, the September 1969 glitch is clearly a resolved, rare event which we have consequently excised from our data sample.[1] To do otherwise would invite possible contamination of the calculated power spectrum if the remaining noise actually stems from a distinct process. On the other hand, if the remaining noise is simply a superposition of many, small, unresolved "glitches" of the same nature as the macro-glitch, then deleting the rare, large event will not effect the form of the calculated power spectrum, but only reduce the overall amplitude.

In such simple cases, excision of apparently resolved events is both prudent and relatively harmless. However not all systems invite such a straightforward appraisal. Consider Her X-1 for example. In figure 2a the observed variation in ν during September 1972 is highly suggestive of a spin-up counterpart to the glitch behavior seen in the Crab and Vela pulsars. That is, superposed on the secular increase in pulse frequency there is a rapid spin-down episode in September 1972, with $\Delta\nu/\nu = 2\text{x}10^{-6}$. Furthermore the December 71 - May 72 data is strikingly noise-free as demonstrated by the constancy of second differences (figure 2c) in this interval. These data raise some difficult issues. If this glitch is superposed on some underlying, distinct noise process which might be analysed separately by excising the glitch, then one must confront the apparently non-stationary statistical properties of that noise process, i.e. the apparent contrast between the 5-6 month "quiet" interval and the remainder of the sample. On the other hand, given the limited quantity of data, one could argue that the variation of ν in figure 2a depicts a single, resolved event (or perhaps only a fraction thereof) whose form follows deterministically (e.g. from the torques between the crust and a lagging superfluid core), and therefore has no statistical significance whatever.

As pure cases, these two apparently divergent interpretations of the Her X-1 data are in a sense degenerate with respect to spectral analysis. The fact is that the spectrum of a single (complete) resolved event has the same expectation as a superposition of many such (unresolved) events. This pleasant simplification allows us to sidestep the issue of resolved versus unresolved events, but leaves the problem of whether or not a complete event has in fact been sampled. Obviously, excision of data (such as the September frequency jump in figure 2a) is potentially disastrous in the case of a single, resolved event. These various ambiguities posed by the Her X-1 data cannot be clarified with the data in hand. We present this example primarily as an indication of how difficult the interpretation of timing observations can be, and that the "close your eyes and turn the crank" brand of objectivity is generally an invitation to err. By contrast, we find that the Crab data are reasonably consistent with the simple case of noise resulting from the superposition of unresolved events, punctuated by rare, resolved events which we now know are of quite a different character.

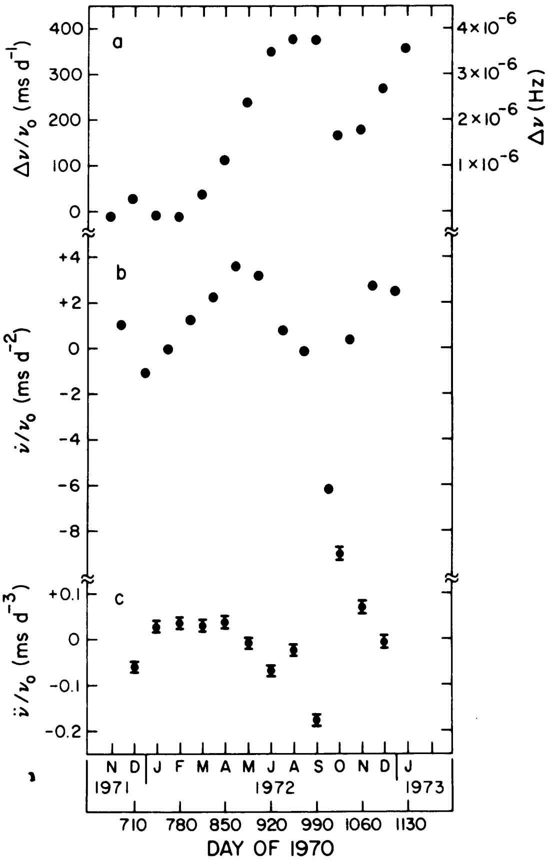

Figure 2: Frequency, $\dot\nu$ and $\ddot\nu$ behavior of Her X-1 over the span of available Uhuru observations. Note significant glitch in September 1972, and possible "pre-glitch" marked by very smooth (low noise), constant $\ddot\nu$ behavior commencing in January of that same year.

In the previous paragraphs we have casually shifted back and forth between descriptions of pulse timing noise variously based on power spectra

https://doi.org/10.1017/S0074180900093037 Published online by Cambridge University Press

of fluctuations in three observables: pulse phase, $P_\phi(\omega)$; pulse frequency, $P_\nu(\omega)$; and pulse angular acceleration, $P_{\dot\nu}(\omega)$. Actually the relationship between these spectra is quite simple. Loosely speaking, the power density at ω in an infinitesimal bandwidth $d\omega$ is proportional to the square of the Fourier amplitude: $[\, A_\omega \sin \omega t + B_\omega \cos \omega t \,]^2$. Since ϕ, ν and $\dot\nu$ are related simply by differentiation, so are the relevant Fourier amplitudes. Thus the spectra are related by successive factors of ω^2 : $P_\nu(\omega) \sim \omega^2 P_\phi(\omega)$, $P_{\dot\nu}(\omega) \sim \omega^4 P_\phi(\omega)$ or $P_\phi(\omega) \sim \omega^{-4} P_{\dot\nu}(\omega)$, etc. In fact, the simplest noise models are the iterated integrals of white noise. Thus if $P_{\dot\nu}(\omega)$ is white (constant power density), then $P_\nu(\omega) \sim \omega^{-2}$ (random walk in frequency) and $P_\phi(\omega) \sim \omega^{-4}$ ("double" random walk in phase). We refer to these latter two forms as 1st-order red noise in frequency and 2nd-order red noise in phase, respectively.[2]

It turns out that concern about the interpretation or even validity of a power spectral analysis of pulse timing residuals is only one of several difficulties which must be confronted in completing such a study. For instance, serious problems are encountered just in trying to compute power spectra of red noise processes. These difficulties are reviewed in detail by Deeter and Boynton (1980), and can be partly indicated by attempting to calculate a spectrum as the square of the Fourier transform of an observed time series. For any finite data set, this local power density estimator falls off as ω^2 toward low frequencies. Clearly, if the spectrum to be recovered is any redder than ω^{-2} , a devastating power leakage problem will occur, rendering this estimation technique invalid. Deeter and Boynton (1980) pose a general mathematical framework leading to a power density estimation technique valid for red, piece-wise power-law spectra calculated from irregularly sampled data. Their modest, one-octave frequency resolution is quite adequate to identify "local features" of the kind associated with the "step" behavior shown in figure 1. Moreover, the calculation of a true power spectrum allows one to actually test for power-law behavior over the available frequency range rather than assume this result as Groth (1975) and Cordes (1980) have done. We emphasize that previously published studies of pulse timing noise could not distinguish between rigid- and nonrigid-body behavior, nor verify the assumption of power-law noise spectra.

A second concern regards limitations on the total noise bandwith which can be examined by this method. The lowest accessible frequency is of course specified by the total length of the data sample. The highest accessible frequency depends on the nature of the intrinsic noise process exhibited by the source through pulse timing residuals. This intrinsic noise is to be distinguished from variability introduced through the measurement process (primarily by counting statistics), which we refer to as "observational noise". The latter contribution is necessarily independent from sample to sample and thus produces white noise in pulse phase measurements. However, a casual inspection of figure 2a reveals that the noise in frequency fluctuations is significantly correlated from one sample to the next, indicating the dominance of power at lower frequencies, or a "red" intrinsic noise process. The pattern of phase fluctuations (the time integral of figure 2a), must be still smoother and thus correspond to

https://doi.org/10.1017/S0074180900093037 Published online by Cambridge University Press

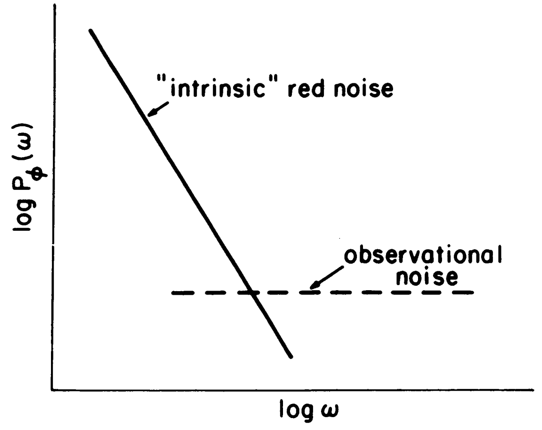

Figure 3: Limitation on observability of intrinsic red phase noise posed by pulse phase measurement uncertainty (observational noise)

an even redder spectrum. This situation is portrayed in figure 3. Here, intrinsic red noise in phase is represented as a power law, and observational noise in phase is shown with power density which is essentially constant. Thus for red spectra, the observability of the intrinsic noise process is generally restricted at high frequencies by observational noise, not by the Nyquist sampling limit. This restriction is discussed in detail by Boynton and Deeter (1979).

V. DISCUSSION

In figure 4 we present spectra of timing fluctuations for two sources, the Crab pulsar and Her X-1. These estimates of $P_{\dot{\nu}}(\omega)$ for both objects are seen to be consistent with a flat (white) spectrum up to the limit imposed by observational noise (dashed line with slope 4). The Crab spectrum is calculated from data of Groth (1975), whereas Uhuru observations provide the base for the Her X-1 analysis.

Since there are no obvious signs here of the kind of structure portrayed in figure 1, rigid body behavior is indicated over the entire range of frequencies explored. Therefore the spectrum is also indicative of the nature of applied torque noise. The fact that both accreting and non-accreting, rotating neutron star systems exhibit approximately white torque noise may be a significant observation; but more striking is the absence of the two-component signature in the power spectrum of the Crab. The model parameters Q and τ_d for the Crab are well-determined from observations of macro-glitch behavior (Anderson et al. 1978). These values, ($Q > 0.9$, $4^d < \tau_d < 15^d$) would produce an unmistakable feature in figure 4, with $P(\text{high})/P(\text{low}) = (1-Q)^{-2} = 600$, lying well within the observable frequency range at $(4^d)^{-1} < \omega_d < (15^d)^{-1}$.

https://doi.org/10.1017/S0074180900093037 Published online by Cambridge University Press

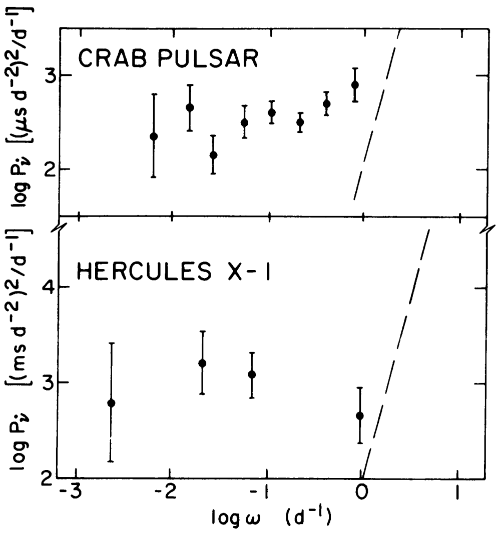

Figure 4: Power spectral density of fluctuations in $\dot\nu$ computed from pulse timing phase residuals for both the Crab and Her X-1. Observational noise contribution is shown by the dashed line.

In figure 5a, we quantify the obvious though the simple hypothesis test suggested in section II. By computing the X^2 surface which results from a relaxation of equation 1 onto the Crab power spectrum, we have constructed the family of X^2 contours corresponding to the labeled significance levels. Curiously, precisely that portion of parameter space occupied by previous estimates of Q and τ_d (marked as points A and B in figure 5a) is most significantly rejected by this test. The implication seems inescapable. Although two-component behavior appears to dominate th dynamical response of the neutron star for sufficiently large amplitude torque perturbations (macro-glitches), the unresolved "micro-glitches" (small amplitude torque perturbations which may arise from a process distinct from that responsible for macro-glitch behavior) do not excite a coupled crust-superfluid response, but instead yield simple rigid body behavior. A similar analysis of the Her X-1 data yields the contours show in figure 5b. Here we have no a priori expectations for Q and τ_d , but can infer that large values of Q are rejected for τ_d between 2 and 30 days.

This demonstration of the failure of the decade-old two-component model of neutron star structure comes as no surprise as it has always beer regarded as the simplest possible description of what is likely a complex system. In fact an extension of this idea, the "three component" model, which includes a component of pinned superfluid vortices, has already bee proposed (Anderson, Pines, Ruderman, and Shaham 1978). If the observed macro-glitch behavior is related to a dynamical interaction of crust and pinned superfluid through the cooperative unpinning process discussed by

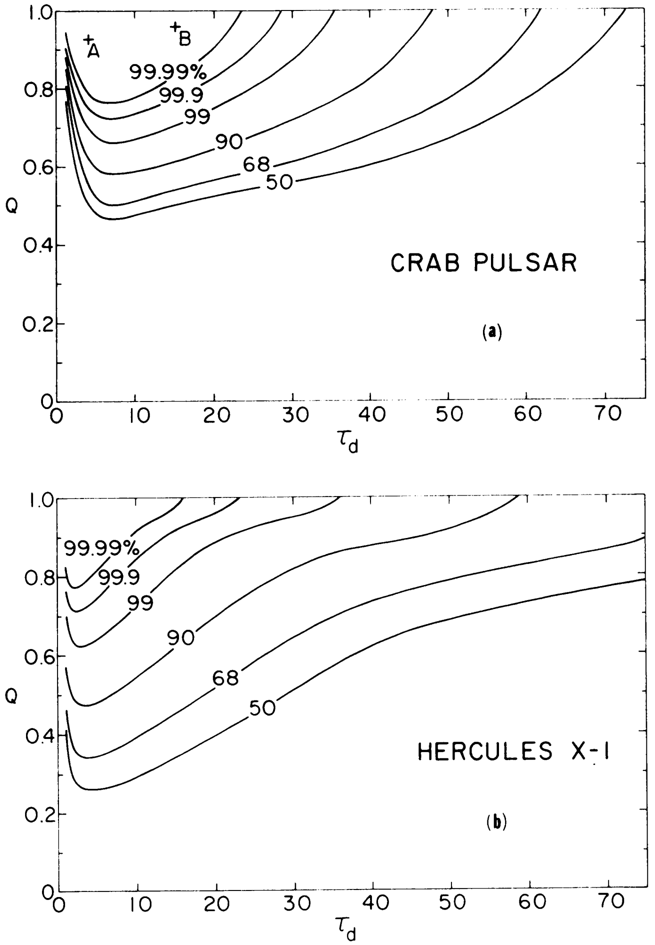

Figure 5: Chi-square surface resulting from a comparison of the "structur function" (figure 1) with the estimated power density (figure 4) for the two sources considered here. Contours are labeled with single-tailed significance level. For the Crab, points A and B label Q, τ_d values compute from observed macro-glitches.

Alpar (see these proceedings), then the "three component" model may alreac contain a natural explanation for the distinction between macro-glitch anc micro-glitch behavior; namely, the existance of an energy threshold for tl initiation of cooperative unpinning which then ensures rigid body behavio: for sufficiently small amplitude torque perturbations. In this case, figures 5a and 5b then restrict the values of Q and τ_d which describe coupling of the crust to the unpinned superfluid component. As already mentioned for the Hercules case, although large values of Q are rejected for $2^d < \tau_d < 30^d$, $Q \leq 0.5$ is allowed on all timescales for both sources.

In concluding, we reiterate that a statistical analysis of pulse timing residuals does not proceed without confronting a host of questions regarding first the appropriateness of a simple noise interpretation, and second, the fundamental difficulties in correctly recovering red noise spectra. Despite these problems, our analysis indicates that the Crab pulsar behavior is well-suited to a noise representation, and is shown to be consistent with a fairly clean power-law spectrum over seven octaves (displaying none of the expected earmarks of internal structure). For Her X-1, although a noise interpretation of spin-rate variations on timescales of months may be appropriate, the form of the variation is highly suggestive of infrequent crust-superfluid relaxation episodes connected by intervals of smoothly building stress (a fully "resolved" process). We are studying evidence for similar behavior in the spinning-up sources Cen X-3 and Vela X-1. Perhaps just those differences which separate the pulsars from the X-ray pulsators (differences in the magnitude and in the sign of angular acceleration, and in the internal temperature), may be exploited to provide the complimentary views of internal structure needed to develope a fuller understanding of neutron star interiors.

Many thanks to J. Deeter, R. S. Boynton, and R. Salvino for their contribution to this study, and to F. K. Lamb, M. A. Alpar and D. Pines for numerous insights and general encouragement. This work was supported in part by the NSF and the Alfred P. Sloan Foundation.

Footnotes:

[1]We use the term "resolved" to describe identifiable, non-overlapping events.

[2]In doing so, we purposely avoid the ambiguity presented by the terms "phase noise" and "frequency noise" when used out of immediate context.

REFERENCES

Anderson, P., Pines, D., Ruderman, M., and Shaham, J.: 1978, J. Low Temp. Phys. 30, p. 839.
Baym, G., Pethick, C., Pines, D., and Ruderman, M.: 1969, Nature 224, p. 872.
Boynton, P. and Deeter, J.: 1979, Compact Galactic X-ray Sources, Eds. F.K. Lamb and D. Pines (UIUC Physics Dept.), p. 168.
Cordes, J.: 1980, Astrophys. J. 237, p. 216.
Deeter, J. and Boynton, P.: 1980, Astrophys. J., submitted.
Groth, E.J.: 1975, Astrophys. J. Suppl. 29, p. 431.
Lamb, F.K.: 1979, Proc. HEAO Science Symp. (NASA CP2113), p. 223.
Lamb, F.K., Pines, D., and Shaham, J.: 1978, Astrophys. J. 224, p. 969.

DISCUSSION

NELSON: It seems exceedingly dangerous to delete sections of data because they "look" singular and distinct from the bulk of the remaining data. What effect on the noise power spectrum of Her X-1 did you get by removing the "stationary region" and the "large torque singularity"?

BOYNTON: As I commented at the time, such a procedure is risky indeed. But non-stationarity cannot be dealt with by excision and we did not delete the "quiet" region you mention. We only meant to illustrate the caution that must be observed in interpreting power spectra in the face of possible non-stationarity. On the other hand, large amplitude, time-resolved features, if distinct from the population of unresolved events (the noise process under study), should be removed since they constitute a superposed process. The question is whether or not they are truly distinct, as you point out. In fact for the feature in the Her X-1 data, which I showed to illustrate this problem, the distinction is not so clear, and the power spectrum is not significantly altered between exclusion and inclusion. These various problems which cloud the interpretation of the Her X-1 spectrum are not apparent in the Crab data (except for the necessary excision of the September 1969 glitch) and we feel confident in our conclusions regarding internal structure of the Crab pulsar.

KUNDT: What are the sizes and repetition rates of the microglitches of the Crab pulsar?

NELSON: As I recall, over a three year period, we have seen on the order of twelve "events" that by their shape appear discrete with amplitudes up to two orders of magnitude below the 1969 glitch. Since the observations are optical we do not have 36 continuous months of coverage.

RUDERMAN: Can you distinguish between positive and negative torque impulses with the analysis you have used?

BOYNTON: I think the only way to determine the sign of events is to actually time-resolve them.

CORDES: Timing noise from ten radio pulsars and the Crab optical data are consistent with events of both signs occurring.

NELSON: In the Crab optical timing data we appear to see small "glitches" of both signs.

ARONS (in response to a question from the audience): "Random walking" of phase residuals is normally regarded as some response of the star to a sort of "noisy" torque, with a specified (assumed simple) power spectrum. Torque really means magnetospheric torque or moment of inertia; therefore even if the magnetosphere is completely quiet, noise can appear from noise in the moment of inertia. In accreting sources, there are plenty of reasons to think the magnetosphere is <u>not</u> quiet, with noisiness of some sort in the torque. Whether the noise is white or red is a question quite a bit beyond the predictive ability of any theory of magnetospheric instability now. In isolated pulsars, there is much less known about possible magnetospheric instability in general, but again, there are lots of reasons for strong fluctuations in the conduction currents assumed in some emission models, and these in turn suggest noisiness in the torque is at least possible.