Review of Scientific Instruments

A flexible data acquisition system for timing pulsars

Cite as: Review of Scientific Instruments **63**, 3551 (1992); https://doi.org/10.1063/1.1143763
Submitted: 10 February 1992 • Accepted: 09 March 1992 • Published Online: 04 June 1998

D. R. Stinebring, V. M. Kaspi, D. J. Nice, et al.

A flexible data acquisition system for timing pulsars

D. R. Stinebring,[a)] V. M. Kaspi, D. J. Nice, M. F. Ryba,[b)] J. H. Taylor, and S. E. Thorsett[c)]
Joseph Henry Laboratories and Physics Department, Princeton University, Princeton, New Jersey 08544

T. H. Hankins
New Mexico Institute of Mining and Technology, Socorro, New Mexico 87801

(Received 10 February 1992; accepted for publication 9 March 1992)

We describe a flexible, inexpensive data acquisition system built for high-precision timing observations of pulsars. The system is designed to interface with a wide variety of radio telescope receiver back ends; it permits standardized measurement techniques and data formats in work carried out at a number of different observatories. Copies of the basic "Mark III" system are now in regular use at the Arecibo Observatory, Green Bank, and the Very Large Array. We describe the specifications, hardware, and software implementation of the system, and briefly outline some of its current applications.

I. INTRODUCTION

Pulsar signals reach the Earth as periodic bursts of broadband radio noise, usually buried in a stronger background of random noise from the Galaxy. The pulse repetition intervals, identified with the rotation periods of spinning, magnetized neutron stars, vary between approximately 1.5 ms and 5 s, while the pulse durations are typically 1%–5% of the period. Over the past 20 years, high-precision timing measurements of these remarkable natural signals have become increasingly recognized as important experimental tools in such diverse areas as general relativity, neutron- and binary-star physics, astrometry, cosmology, and the physics of the interstellar medium. In most of these applications, the interesting science requires the highest possible measurement precision. Therefore, data acquisition systems built especially for accurate pulsar timing, with special attention to details of the apparatus involving time and frequency control, have become highly desirable equipment at radio observatories.

The best pulsar timing observations have been made by sampling the detected signals 100 or more times per pulsar period, synchronously adding samples collected at the same pulsar phase (referred to as "phase bins"). Thus a single observation results in a single average periodic waveform ("pulse profile") for each data channel. The location of the pulse within the averaging window, together with the recorded start time of the first sample, determines the desired topocentric pulse time of arrival. In well-executed measurements, the ultimate timing accuracy is proportional to the observational signal-to-noise ratio. Other parameters being equal, sensitivity and timing accuracy are enhanced in proportion to $B^{1/2}$, where B is total received bandwidth; however, propagation of the received signals through the dispersive interstellar medium broadens the apparent pulse widths observed in a wide receiver passband, degrading the accuracy of times of arrival. Therefore, to make effective use of a large total bandwidth, many narrow receiver channels must be employed and sampled simultaneously. In practice, banks of several dozen or more spectral channels are typically used. When the shortest-period millisecond pulsars are being observed, the overall data rate can easily reach several megasamples per second.

Synchronous signal averaging over intervals of a few minutes reduces the data rate to manageable levels, but in order to maintain high precision the averaging must be done with careful attention to detail. Apparent pulsar periods change with time in a predictable manner, owing to motions of the observatory and the pulsar (as well as the intrinsic pulsar spin-down). Consequently, it is necessary to use a dynamically calculated period to synchronize the signal averaging hardware. In the important case of binary pulsars with fast, highly eccentric orbits (such as PSR 1913+16, with pulsar period $P=59$ ms, orbital period $P_b=7.8$ h, and orbital eccentricity $e=0.617$), the synchronizing frequency must be adjusted as often as once per second in a phase-continuous manner. Accurate time stamps for the integrations are also required. The effective "arrival time" of a series of integrated profiles can, in the best cases, be measured to tenths of a microsecond or better—so the ultimate time reference, as well as any intermediate steps in time calibration, should be at least this accurate.

The Princeton Mark III pulsar timing system was developed with these requirements in mind. In minimal form, it consists of a PC/AT-style computer with three special-purpose cards plugged into the backplane. The special hardware, also connected to the receiver and to various time and frequency signals, provides rapid sampling and A/D conversion of many incoming data channels, as well as synchronous signal averaging. The averaged data are transferred to computer memory every few minutes, and stored on disk for later analysis. The pulse profiles obtained during each integration (including individual spectral channels and a "dedispersed" total) are displayed graphi-

---

[a)]Present address: Department of Physics, Oberlin College, Oberlin, OH 44074.
[b)]Present address: MIT Lincoln Laboratory, Lexington, MA 02173.
[c)]Present address: Owens Valley Radio Observatory, 105-24, California Institute of Technology, Pasadena, CA 91125.

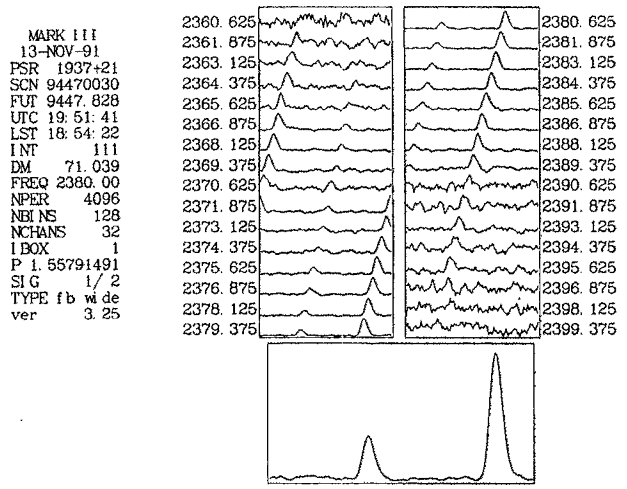

FIG. 1. Rendition of a typical Mark III screen display. Data are from PSR 1937+21 and were collected at Arecibo Observatory. Text information in the display describes the current system setup as well as providing incidental information such as data and time. The 32 small plots are integrated pulse profiles from individual receiver channels. (Interstellar propagation effects cause dispersive delay of the pulsed signal at lower frequencies and strong amplitude scintillations that vary with frequency and time.) The larger plot is a "de-dispersed" profile in which the 32 channels have been aligned and summed.

cally to provide a rapid check of system integrity at the time of observation. Figure 1 illustrates a sample display obtained during an observation of millisecond pulsar PSR 1937+21 at the Arecibo Observatory.

Flexibility in the basic design has allowed us to implement variations of the Mark III system which readily accommodate the requirements of different observations—for example, taking advantage of diverse receiver "back ends" that are locally available. We have used the system only at radio observatories, but it could easily be adapted to work with other suitable analog signals, such as the output of a photomultiplier tube for optical pulsar observatories. The overall design has proven reliable, and maintenance, when required, is fast and easy. Modular circuit boards allow for quick replacements "in the field," and actual repairs can be made later, in the laboratory.

II. HARDWARE

The basic Mark III hardware includes three distinct types of PC-compatible circuit board: a multiplexer and A/D converter, a signal averager, and a programmable frequency synthesizer with phase-continuous output. Specifications of the generic system are summarized in Table I, and a block diagram is presented in Fig. 2. In practice we often use two or more multiplexer/signal-average pairs to record data simultaneously from several receiving systems. A single synthesizer produces all necessary synchronizing signals. Different construction techniques have been used for the three boards, as appropriate to the circuit complexity and necessary density of integrated circuits. Our production versions of the frequency synthesizer and signal averager were commercially wire wrapped, while the multiplexers are built on relatively simple two-layer printed circuit boards.

TABLE I. Specifications of the Mark III data acquisition system.

Multiplexer A/D:	
Input channels	2, 4, 8, 16, 32
Maximum sample rate	2.8 MHz
Sample precision	6 bits
Signal averager:	
Data memory	32k×24 bits
Valid sample memory	32k×16 bits
Phase bins per period	⩽ 32 768/(numbers of channels)
Boxcar integration factor	1–255
Periods integrated	1–65 535
External synchronization accuracy	±25 ns
Frequency synthesizer:	
Maximum master clock frequency	20 MHz
Output frequency	0–5 MHz
Resolution	0.0046 Hz

The Mark III system has a simple interface to other relevant observatory hardware. Analog signals, in the form of up to 32 independent voltages per multiplexer/signal-averager pair, are introduced via two 40-conductor ribbon cables plugged into the multiplexer board. An external frequency standard (running at any frequency up to 20 MHz) provides the synchronous clock for the entire system, and an "external tick" (typically 1 or 6 pulses per minute) is used to synchronize all hardware and time stamp the data. Both these signals are derived from an external reference atomic clock. Additional communication is available through standard parallel and serial ports on the PC, allowing information exchange with a telescope control computer for automated observing, and for other external signals that may be needed.

Internal communication between system components

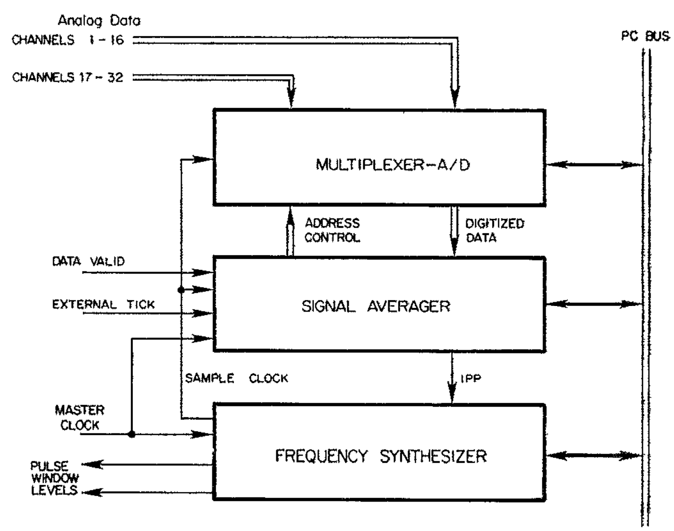

FIG. 2. Block diagram of the Mark III data acquisition system. Several multiplexer-A/D and signal-averager pairs can be used in a single system to accommodate more input channels. The "pulse window levels" can be used to synchronize external equipment with the pulsar period.

takes place over dedicated ribbon cables between the circuit boards. The PC bus is used for computer control of a number of programmable registers and for transfer of data between the signal averager and main computer memory. The computer bus is not used for internal Mark III signals, however. During data acquisition, the system is nearly free running and requires little interaction with the computer, which is therefore free to look after other details such as displaying results from the previous integration, communicating with the operator, and storing data on disk.

A. Multiplexer A/D

Detected signals from the telescope enter the Mark III system through one or more 32-channel analog multiplexer-A/D boards.[1] These units can be clocked at rates up to 2.8 megasamples per second, and the input signals are digitized and transferred to the signal averager card at the same aggregate throughput rate. Each of 32 receiver channels can therefore be sampled at intervals as short as 12.8 μs. A smaller number of channels can be sampled proportionally faster. To achieve this performance at modest cost, we implemented the circuit with two identical 16-channel systems, each using a Harris H-516 analog multiplexer, two stages of signal conditioning in LM318 op-amps, and CA3306 six-bit flash A/D converter. Channel addressing and sampling alternates between the two 16-channel circuits in an interleaved fashion, thereby providing adequate settling times (700 ns at the fastest usable sample rates) for the multiplexed signals. Toggling between the two 16-channel portions of the circuit is controlled by the least significant bit of the signal-averager address lines (see Sec. II C), which enables the appropriate A/D outputs during a particular cycle of the sample clock. The next four address bits control the addressing of the 16-to-1 multiplexers—in one case with a one clock-cycle delay, provided by latch registers in a programmable array logic (PAL) circuit. Bit masking done in the PAL chip, with a mask word sent from the computer, also provides for programmable selection of 2, 4, 8, 16, or all 32 signal channels.

B. Signal averager

The signal-averager board[2] contains the memory in which the digitized data are accumulated, along with appropriate control circuitry. The data memory is composed of three 32k$\times$8-bit static random-access memory chips, addressed in parallel to form 32k words of 24 bits each. Each word accumulates the sum over many periods of digitized data samples for one particular combination of phase bin and data channel. A second set of memory, 32k$\times$16 bits, counts the number of samples added into each word of data memory. In general this count would be the same for all memory addresses; however, an external "data valid" signal line can be used to flag incoming data which is invalid and therefore should not be added into the accumulated sum.

The signal averager is initialized by loading several register values: the number of data channels, number of phase bins per pulsar period, number of adjacent samples to sum (boxcar integration[3]), number of periods to integrate, and two numbers describing location and length of a window within the period during which data are to be collected. Each of these registers occupies a distinct address in the PC I/O space. In order to allow communication with one or all of multiple signal averagers, each one has an additional "select" register with a unique address (set by on-board switches). Several signal averagers can be simultaneously in the "selected" state, and all of them will respond in parallel to most commands. Obviously, data readout from multiple signal averagers must be done sequentially.

Before data collection begins, a once-per-period (IPP, for "interpulse period") synchronizing signal is initialized by starting a counter a fixed length of time after an external tick. The size of this initial delay is set in a programmable register, and is measured by counting master clock cycles. Once started, the IPP signal is continuously generated by counting sample-clock cycles (see below), and a counter overflow triggers the start of the next period. The IPP signal is passed to the frequency synthesizer, which uses it to control external "pulse window" lines synchronously with the pulse period.

Data collection is begun by sending a command over the PC bus. Once the start command has been received, the circuit waits for an external tick and starts actual data collection at the next IPP signal after this tick. Master clock cycles between the external tick and the start of data collection are counted. Since both the master clock and the start tick are generated from very precise external reference signals, the data collection start time is known to high accuracy, approximately half a master clock period.

Once data collection commences, a single sample is collected from some data channel during each cycle of the sample clock (generated by the frequency synthesizer as described below.) Selection of the data channel and the memory location into which the sample is added is controlled by a set of three cascaded counters, which together manipulate memory and multiplexer addressing to properly accumulate the samples. An n-bit counter cycles through all 2^n input channels, selecting each one in turn. The n-bit counter ($2 \leqslant 2^n \leqslant 32$) is used both as the low-order part of the memory address and as the address sent to the multiplexer-A/D board to select the appropriate data channel. At the end of each pass through the 2^n data channels, the boxcar integration counter is incremented, and the cycle repeated. Once the appropriate number of samples have been boxcar integrated (perhaps only one), the phase bin counter is incremented. This counter is used for high-order addressing of the memory. No interaction between the PC and the card takes place during data collection.

Data readout is accomplished in a separate mode of operation, and cannot be done in parallel with data collection. During the read operation, memory addresses are cycled as above, with the memory address controllers incrementing as soon at the PC bus indicates the PC has latched the data. The low-order words of the data, the high-order bytes of the data, and the valid-sample-count words are successively made available to the PC. Depending on the computer speed and quantity of data collected, it typically takes a few seconds to extract the full contents of signal-averager memory and initiate a new integration.

The signal-averager circuitry relies heavily on PAL circuits for such things as decoding PC bus signals, controlling memory addressing, and interpreting clock ticks. The rapid manipulation of memory I/O during data accumulation requires a complex sequence of control signals generated at a fast rate, up to 20 MHz. We derive these from a PROM-based algorithmic state machine, implemented with two parallel TBP38S030-15 chips.

C. Frequency synthesizer

Changing pulsar periods due to Doppler shifts caused by the Earth's and pulsar's relative motions require the frequency of data sampling to be updated regularly in order to maintain a fixed integral number of samples per pulsar period. The frequency changes must take place with no missing or extra pulses in the sample clock. The Mark III system uses a programmable synthesizer[4,5] based on the "direct synthesis" technique. On each cycle of the master clock, a digital phase increment is added to a 32-bit accumulator, and an output pulse is generated whenever the accumulator overflows. The phase-increment register can be reloaded under computer control at any time; in practice, our standard software does this once each second. With a 20-MHz master clock, the synthesizer provides a useful frequency range 0–5 MHz, with 0.0046-Hz resolution. An internal crystal oscillator is available for system tests, although in actual use an external frequency standard is used for the master clock in order to provide the necessary 10^{-10} fractional accuracy and short-term stability for pulsar timing.

The synthesizer's principal function in the Mark III system is to produce a series of pulses at the desired data sampling rate. These pulses become the sample clock used in the signal averager, controlling all data-acquisition timing by means of the sequencing of multiplexer and memory addresses described above. The required sample-clock frequency is the product of topocentric pulsar frequency, the number of sampled receiver channels, the number of phase bins per pulsar period, and the integer boxcar-integration factor. All of these parameters are user selectable under software control at the start of an integration.

Nine programmable interval timers are also provided on the frequency synthesizer board. Three of these generate 1-ms, 1-s, and 10-s ticks that can be synchronized with the external reference tick or used in place of it, for testing purposes. One timer generates the sample clock by counting the synthesizer output. Four timers count the sample clock and generate buffered, pulsar-period-synchronous logic signals which can control external equipment such as a pulsed noise source, for receiver gain calibration. These "pulse window levels" can be turned on and off at selectable pulse phases.

III. SOFTWARE

The host computer for the Mark III system runs under the MS-DOS operating system, with an added real-time executive providing rudimentary multitasking ability. Most of the code is written in Fortran; the real-time extensions, written in 80x86 assembly language, are linked with other compiled object modules to produce a single executable file. When the Mark III software is active, the multitasking features permit synchronous updating of the frequency synthesizer, concurrent with asynchronous tasks such as keyboard input, updating the display screen, disk I/O, and associated calculations. When the Mark III system software terminates, the real-time executive exits gracefully and the computer behaves as a normal MS-DOS machine.

The real-time code executes in response to the 18.2-Hz clock-interrupt signal present in all IBM PC-compatible computers. Its primary purpose is to determine the exact time, by reading a counter on the synthesizer board which counts the master clock and is reset by the external reference tick. The interrupt routine keeps track of counter resets by incrementing suitable variables in memory. At 1-s intervals, with a maximum phase jitter of 27 ms, the interrupter routine sends updated frequency requests to the synthesizer. It then computes the frequency that will be needed 1 s hence, by evaluating a time polynomial whose coefficients are stored in a common memory block. (The necessary calculations of topocentric pulsar phases, stored compactly in the form of polynomial coefficients, are carried out before an observing session begins, by a separate Fortran program.) The interrupt routine completes its work within a few milliseconds; on exit it fully restores the state of the computer as seen by the remainder of the program, including the contents of the 80x87 floating-point coprocessor.

A main program attends to all remaining tasks required for pulsar observations. Upon start-up, it reads a set of configuration files which describe the observing site and the number and identification of all boards in the system. The PC clock interrupt vector is reset to the entry point of the real-time interrupt routine. The presence of external clock signals is verified, and the PC software clock is synchronized with the reference tick. The program then initializes all Mark III hardware and awaits further instructions.

Observing sequences are started either in immediate response to keyboard requests or according to a list of times and pulsar names read from a file for automated observing. The program selects polynomial coefficients to be used by the interrupt routine, stores them in common memory, and starts the signal averager. While an integration is in progress, the display screen is continually refreshed with alphanumeric information and the keyboard is monitored for any additional commands. The signal averager runs asynchronously, without attention from the computer. When a "done" bit signifies an integration is finished, data are moved from signal-average memory to computer memory, and then to disk, along with essential book-keeping information such as date and time, pulsar name, observing frequency, and other relevant details. A new scan is started once the data have been read into computer memory. The computer plots the single-channel and dedispersed profiles on the screen for operator inspection.

IV. APPLICATIONS AND SUMMARY

The Mark III data acquisition system can be duplicated for approximately $3000, transported easily, and used with a wide range of observatory hardware. A nearly final version was first tested in late 1987 at the Arecibo Observatory. Since then, additional systems have been installed at the NRAO facilities in Green Bank, WV and Socorro, NM (the VLA). By now all three systems have been used by our group and many others to address a wide range of scientific questions.

The Arecibo implementation includes a frequency synthesizer and three signal averager/multiplexer pairs, as well as auxiliary circuitry that carries out coherent (predetection) dispersion removal[6–10] and Stokes-parameter polarimetry.[11] A noteworthy added feature of this system is its ability to compute scintillation spectra at the end of each integration, automatically retuning the narrow bandwidth de-dispersers to take advantage of signal maxima (see Fig. 1). The equipment sees regular use in several long-term projects for monitoring the pulse arrival times of millisecond and binary pulsars, with particular scientific emphasis on goals involving tests of gravity,[12–15] tests of cosmological models,[16] and detection of planetary companions of pulsars.[17] It has also been used for measuring the detailed polarization properties of millisecond pulsars,[11,18] determining the component masses in the PSR 1855+09 binary pulsar system,[19] and exploring bizarre eclipse phenomena in the millisecond pulsar PSR 1957+20.[20–22]

Since early 1989, a Mark III system has made almost daily timing and flux-density measurements of 35 strong pulsars with one of the NRAO 85-foot telescopes in Green Bank.[23] Such observations are important for study of the intrinsic timing noise found in certain young pulsars, and the regular flux-density measurements provide uniquely valuable data for studying inhomogeneities in the interstellar medium.[24] The Green Bank installation uses the automatic schedule-following capabilities built into the Mark III software, and runs for weeks at a time with little operator intervention. Some further software extensions and intercomputer communications allow this system to reduce its own data automatically, sending the results to us by electronic mail.

A Mark III system was installed at the VLA in 1990, and is now being used for observations of a variety of millisecond and binary pulsars with particularly interesting characteristics. One completed application involves an extensive study of the eclipsing binary pulsar PSR 1744−24A.[25–27] We have also begun a series of timing observations of a number of important pulsars outside the declination range of the Arecibo telescope;[11] this work is expected to continue for some time.

ACKNOWLEDGMENTS

We thank G. Berman, G. Bower, M. Convery, W. Dunham, J. Klein, M. McKinnon, V. Polyakov, A. Tollsten, A. Vázquez, and J. Yamamoto for contributions to various phases of development of the Mark III system. Funding has been provided to the project by the National Institute of Standards and Technology, the National Science Foundation, and the Dudley Observatory, and we gratefully acknowledge graduate fellowships from the General Electric Foundation, the National Science and Engineering Research Council of Canada, the U. S. Department of Education, and the NSF.

[1] G. Berman, Princeton University Physics Department internal report, 1989.

[2] W. J. Dunham, M. E. thesis, Thayer School of Engineering, Dartmouth College, 1987.

[3] Boxcar integration permits varying the effective integration time without changing the post-detection low-pass filters in each receiver channel, thereby allowing a nearly constant fractional-period resolution over a wide range of pulsar periods.

[4] A. Tollsten, Senior thesis, Thayer School of Engineering, Dartmouth College, 1987.

[5] J. Tierney, C. M. Rader, and B. Gold, IEEE Trans. Audio Electroacoust. **AU-19**, 48 (1971).

[6] T. H. Hankins and B. J. Rickett, Meth. Comp. Phys. **14**, 55 (1975).

[7] J. Rajkowski, M. E. thesis, Thayer School of Engineering, Dartmouth College, 1987.

[8] T. H. Hankins and J. M. Rajkowski, Rev. Sci. Instrum. **58**, 674 (1987).

[9] T. H. Hankins, D. R. Stinebring, and L. A. Rawley, Astrophys. J. **315**, 149 (1987).

[10] M. F. Ryba, Ph.D. thesis, Princeton University, 1991.

[11] S. E. Thorsett, Ph.D. thesis, Princeton University, 1991.

[12] J. H. Taylor and J. M. Weisberg, Astrophys. J. **345**, 434 (1989).

[13] T. Damour and J. H. Taylor, Astrophys. J. **366**, 501 (1991).

[14] J. H. Taylor, A. Wolszczan, T. Damour, and J. M. Weisberg, Nature (London) **355**, 132 (1992).

[15] T. Damour and J. H. Taylor, Phys. Rev. D (in press, 1992).

[16] D. R. Stinebring, M. F. Ryba, J. H. Taylor, and R. W. Romani, Phys. Rev. Lett. **65**, 285 (1990).

[17] A. Wolszczan and D. A. Frail, Nature (London) **355**, 145 (1992).

[18] S. E. Thorsett and D. R. Stinebring, Astrophys. J. **361**, 644 (1990).

[19] M. F. Ryba and J. H. Taylor, Astrophys. J. **371**, 739 (1991).

[20] A. S. Fruchter, D. R. Stinebring, and J. H. Taylor, Nature (London) **333**, 237 (1988).

[21] A. S. Fruchter, G. Berman, G. Bower, M. Convery, W. M. Goss, T. H. Hankins, J. R. Klein, D. J. Nice, M. F. Ryba, D. R. Stinebring, J. H. Taylor, S. E. Thorsett, and J. M. Weisberg, Astrophys. J. **351**, 642 (1990).

[22] M. F. Ryba and J. H. Taylor, Astrophys. J. **380**, 557 (1991).

[23] D. J. Nice, J. H. Taylor, and D. R. Stinebring, Bull. Am. Astron. Soc. **22**, 1286 (1990).

[24] V. M. Kaspi and D. R. Stinebring, Astrophys. J. (in press, 1992).

[25] D. J. Nice, S. E. Thorsett, J. H. Taylor, and A. S. Fruchter, Astrophys. J. Lett. **361**, L61 (1990).

[26] S. E. Thorsett and D. J. Nice, Nature (London) **353**, 731 (1991).

[27] D. J. Nice and S. E. Thorsett, Astrophys. J. (submitted).