

The S2 baseband processing system for phase-coherent pulsar applications

R. Wietfeldt^{1,*}, W. Van Straten², D. Del Rizzo², N. Bartel², W. Cannon^{1,2}, and A. Novikov^{1,2}

¹ Institute for Space and Terrestrial Science, 4850 Keele St., Toronto, Canada, M3J 3K1, Canada

² York University, 4700 Keele St., Toronto, Canada, M3J 1P3, Canada

Received July 1, 1997; accepted March 24, 1998

Abstract. The S2 baseband processing system (S2-BPS) is designed for a variety of pulsar measurements such as ultrafast signal fluctuations and pulse timing. Its key features are: the S2 recorder; the S2 Tape-to-Computer Interface (S2-TCI) which provides an efficient high data rate transfer from S2 tapes to computer; and the tightly coupled control interface to the computer which enables automated data transfer and processing. Developed originally for VLBI applications, the S2 recorder is based on the use of commercial VHS tape transports (VCR's), modified for use in digital high density, high data rate applications. A single S2 recorder "tape-set" of eight SVHS tapes provides up to 500 GBytes of data storage, and an unattended operating time of up to 8.5 hours at the maximum data rate of 128 Mbits/s or 16 MBytes/s, corresponding typically to 16 MHz bandwidth in two circular polarizations at 2-bit quantization, which generates a data rate of up to 1 GByte/minute. With 16 MHz individual record channels, a single S2 yields a sample time resolution of 31 ns. For more challenging applications, multiple S2's may be used for increased bandwidth and higher time resolutions. In this paper, we present an overview of the S2 baseband processing system and plans for future implementations of the system, and demonstrate first results of Vela pulsar recordings at Parkes in September 1995. We pay specific attention to the S2-TCI, a system designed to transfer S2 recorded data from tape to computer memory (RAM) via direct memory access (DMA). Computer data may be processed locally or sent to a more powerful "remote" (super) computer for processing.

Key words: pulsars — radio astronomy — baseband recording — coherent de-dispersion

1. Introduction

Pulsars are relatively weak radio sources with steep spectra, with signals which are highly polarized and variable on time-scales from sub-microseconds to years (for a recent compendium, see Johnson et al. 1996). High time-resolution observations have resulted in a deep insight into the degree of coherence of the pulsar radiation and the physical emission mechanism itself. Most importantly, such observations can be made for sensitive searches for companions of pulsars like planets (Wolszczan 1994) and perhaps black holes, and for the investigation of several issues in fundamental physics. In particular, timing measurements have revealed that millisecond pulsars are extremely stable clocks rivaling in their long-term stability cesium clocks on Earth, on which terrestrial time standards are based (Taylor 1991). Accurate timing of binary pulsars allows studies to be made of the nature of gravity and tests of general relativity (Taylor 1994). When pulsar timing is combined with VLBI measurements, it can be used to tie the solar system reference frame directly to the quasi-inertial extragalactic reference frame for improved determinations of earth orbital parameters, tests of general relativity, and spacecraft navigation (e.g. Bartel et al. 1996). And finally, timing permits searches of the cosmological background of gravity waves predicted as a remnant from the early universe (Stinebring et al. 1990).

There are three phenomena that limit the time resolution obtainable for pulsar signals: interstellar scattering, interstellar dispersion, and the rise-time of the receiving system. Interstellar scattering broadens pulses through multipath propagation. The broadening time, t_{ISS} , is proportional to $\nu^{-4.4}$ for observing frequency ν , and is in general negligible for $\nu > 1.4$ GHz. For instance, for pulsars with a dispersion measure $DM \leq 150$ pc cm⁻³, t_{ISS} is still smaller than the shortest S2 sampling time of 31.25 ns at $\nu \geq 1.4$ GHz (Hankins 1996). In contrast, interstellar dispersion has, in most cases, a more severe effect on the obtainable time resolution. It changes the relative phase

Send offprint requests to: R. Wietfeldt

* Present address: Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., MS 264-832, Pasadena, CA 91109, U.S.A.

of the frequency components of the pulsar signal which is equivalent to smearing of the signal over a time t_{DM} across the filter bandwidth $\Delta\nu$. This dispersion time is given as $t_{DM} = 8.3 \cdot 10^{15} DM \Delta\nu / \nu^3$. The highest time resolution can be obtained by choosing $\Delta\nu$ such that the filter rise-time, $\sim 1/\Delta\nu$, equals t_{DM} . In the case of our pulsar example above, the highest time resolution would be $\sim 20 \mu\text{s}$, obtainable with a filter of 50 kHz bandwidth, chosen appropriately for the combination of pulsar DM and observing frequency. This time resolution is still almost three orders of magnitude coarser than t_{ISS} and the shortest sampling time of the S2. Observations at other frequencies and/or of other pulsars generally require filters with different bandwidths to match t_{DM} with $1/\Delta\nu$. The signal from the filters is typically applied to square-law detectors where the relative phase of the frequency components is inevitably lost. The most effective and flexible way of correcting for the signal phase distortions induced by the interstellar medium (ISM) is by recording the baseband signal and convolving it with the inverse impulse response of the ISM before square-law detection. This way, the obtainable time resolution of the pulsar signal is limited only by the broadening time of the interstellar scattering and the Nyquist sample time of the baseband recording system. This technique was pioneered by Hankins (1971) and used subsequently for particular studies of pulsar radiation (see Hankins 1996 for a review). At the time, only relatively small data sets could be analyzed because of the need for a huge sampling rate and the limitations in storage capacities and computing power. Today these capacities and powers are orders of magnitude larger than those of two decades ago, and a new approach towards phase-coherent pulsar observations and subsequent processing with a general-purpose computer is now feasible.

In this paper, we describe a phase-coherent baseband processing system for radio pulsar observations. Short reports on this system were given earlier by Wietfeldt et al. (1995, 1996a). Subsequently, a report on another new baseband processing system was given by Jenet et al. (1997). Our system is based on the S2 VLBI recorder (Wietfeldt et al. 1996b) developed at the Institute for Space and Terrestrial Science (ISTS) at York University in Toronto, Canada. At present, approximately 20 radio astronomical observatories worldwide are permanently equipped with an S2 recorder. Each of these stations could in principle be utilized for phase-coherent baseband recording of pulsar signals on commercially available SVHS tapes. After the observations, the recorded tapes may be processed with a computer/workstation at particular processing centers. In the remainder of this paper, we give an overview of the S2 baseband system in Sect. 2, discuss considerations for the design of the Tape-to-Computer Interface (S2-TCI) in Sect. 3, describe the software control system of the S2-TCI in Sect. 4, present results of the first observations in Sect. 5, and derive our




Fig. 1. S2 Baseband Processing System (S2-BPS) Overview. Baseband data are recorded on the S2 at up to 128 Mbits/s. S2 tapes are shipped to the processing center, where data are transferred to computer via the S2 Tape-to-Computer Interface (S2-TCI) at up to 16 MBytes/s, with subsequent processing in software. In this figure, only the “local” computer is shown; the optional “remote link” shown provides the mechanism for data transfer to a more powerful computer for more efficient processing

conclusions and present future prospects for baseband processing systems in Sect. 6.

2. S2 baseband processing system overview

An overview of the S2 baseband processing system is shown in schematic form in Fig. 1. The S2 recorder and computer/workstation are standard available items; in particular the S2 recorder is produced and available from ISTS (Cannon et al. 1997). The S2-TCI system configuration is a master/slave configuration with the computer controlling the S2 recorder and data transfer. The data transfer is implemented via direct memory access (DMA), with S2 playback data transferred to computer RAM. Typically, data are either processed immediately from RAM or sent to disk for later processing from disk; alternately, as described later, data may be sent directly to a more capable (super) computer. The hardware interface between S2 and computer is via a custom-designed Computer Interface Card (CIC) residing within the S2, and a commercially-available DMA Card (DMAC) residing within the computer. This interface consists of 16 data lines, a data transfer clock with a DMA transfer rate of 16 MBytes/s or 8 MWords/s where a word is defined as 16 bits, and a “SYNC” signal designating start-of-transfer, all sourced by the S2. The software monitor/control interface between S2 and computer is via the S2 Recorder Control Link (RCL) which uses standard ethernet network communication protocols. The S2-TCI monitor/control program, dubbed *s2tci*, has been designed to provide simple, automated data transfer from S2 recorder to computer.

2.1. Observatory data acquisition system

The function of a radio telescope data acquisition system is to generate digitized baseband voltage signals, representing the incoming electric field, suitable for recording onto videotape. The basic elements are typically:

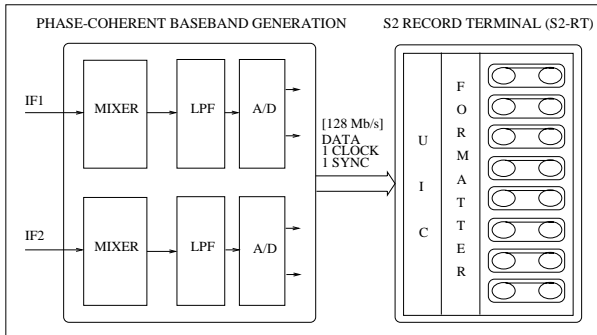


Fig. 2. S2-BPS – Baseband Acquisition System. Pulsar baseband data are recorded on an S2 Record Terminal (S2-RT) at up to 128 Mb/s (16 MBytes/s). For simplicity, two IF input channels and “real”-signal baseband generation at 2-bit quantization are shown. Data samples recorded to S2 tape are time-tagged synchronized to the observatory UTC time-reference “SYNC” (1 Hz) signal

- RF-to-IF down-conversion;
- IF-to-baseband conversion, generating either “real” or “complex” baseband signals;
- Baseband low-pass (anti-aliasing) filtering;
- Digitization, or analog-to-digital (A/D) conversion, consisting of sampling and quantization steps. Sampling of a “real” signal of baseband bandwidth B MHz occurs at the Nyquist rate of $2B$ Msamples/s; sampling of a “complex” signal, i.e. of two signals I and Q, each of bandwidth B MHz, occurs at B Msamples/s for each of I and Q. Quantization: typically 1- or 2-bits per sample is used, although higher-level quantization is often used;
- High-speed synchronization of digitized samples to UTC;
- Digitized baseband signal recording.

These functional steps of a generic baseband acquisition system are displayed in Fig. 2, where the S2 recorder is shown as baseband signal recorder.

2.2. S2 recorder

Developed originally for VLBI applications, the S2 recorder is based on the use of commercial VHS tape transports (VCR’s), modified for use in digital high density, high data rate radio astronomical applications (Wietfeldt et al. 1996b). There are two forms of S2 recorder: an S2 Record Terminal (S2-RT) and an S2 Playback Terminal (S2-PT). These machines are identical except for the number of data Decoders; an RT has only one Decoder for real-time diagnostics during recording, whereas a PT has eight Decoders for full bandwidth data playback capability.

At the radio telescope, astronomical signals from the data acquisition system signal digitizer(s) enter the S2-RT User Interface Card (UIC), which provides the input/record interface hardware at the telescope. The data

are passed from the UIC to the S2 Formatter and to videotape. The primary role of the Formatter is to insert UTC time-stamps onto videotape to identify data samples recovered from tape on playback. The Formatter overhead information, which includes sync words, UTC and auxiliary user information such as observing mode, raises the effective data rate to 17.44 Mb/s/transport to implement the S2 “transparent” or “non data replacement” data format. The advantage of this data format type is that none of the input data are overwritten or lost, a feature particularly important in the pulsar application. At the processing center, an S2-PT reproduces the digitized data exactly, apart from tape errors, as it existed at the radio telescope digitizer(s) output. This is accomplished via eight Decoders (one per transport) in the S2-PT. Each Decoder recovers the recorded data and time-stamps off tape, and passes its 16 Mb/s deformatted data to the UIC, which reproduces the original digitized data and the record timing. The recovered UTC is available over the S2 Recorder Control Link (RCL) every second, synchronous to the S2 1 Hz (“SYNC”) reference marker at the UIC output.

At full bandwidth, a single S2 utilizing eight VHS transports with currently available super-long-play (SLP) thin tape video cassettes records at a maximum aggregate rate of 128 Mb/s (16 MBytes/s) for an unattended operating time of up to 8.5 hours and total data capacity of up to 4 Tbits or 500 GBytes. The unattended operating time may be extended beyond 8.5 hours for narrower recording bandwidths. Numerous channelizations are supported by the S2. Up to 16 digital bit-streams at up to 32 Mb/s/bit-stream may be recorded, such that the 128 Mb/s aggregate data limit is not exceeded. The S2 has been designed as an “intelligent” machine, requiring simple initial configuration and subsequent automated operation. It may be operated remotely using standard ethernet network communication (TCP/IP telnet and socket) protocols.

2.3. S2 tape-to-computer interface (S2-TCI)

The role of the S2 Tape-to-Computer Interface (S2-TCI) is to transfer data recorded on S2 tapes to computer. The S2-TCI is shown schematically in Fig. 3. The S2-TCI consists fundamentally of an S2-PT and a computer¹. Two additional cards are required in addition to the standard S2 and computer systems: a Computer Interface Card (CIC) within the S2; and a DMA Card (DMAC) within the computer. The roles of these cards are described below. The

¹ Currently, Sun Sparc-stations/workstations running SunOS or Solaris operating systems are supported, although other platforms may be supported in the future. In particular, we are considering the development of direct connections to super computers for more efficient data transfer and processing capability.

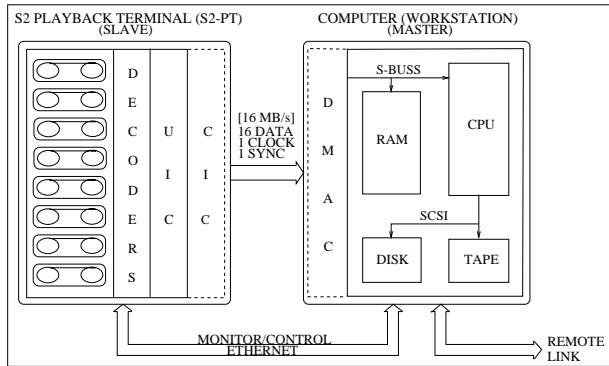


Fig. 3. S2-BPS – Processing Center S2-TCI System Overview. Pulsar baseband data recorded on an S2 Record Terminal (S2-RT) at 128 Mbits/s (16 MBytes/s) are played back on an S2 Playback Terminal (S2-PT). Data are transferred to computer RAM via DMA at up to 16 MBytes/s and to a “local” output device, typically RAM/disk; the optional “remote link” shown provides the mechanism for data transfer to a more powerful “remote” computer for data processing

interface between S2 CIC and computer DMAC consists of 16 data lines, data transfer clock (typically 8 MHz, such that 16-bit data transfers yield the 128 Mbits/s or 16 MBytes/s maximum rate), and a SYNC signal designating start-of-transfer, all sourced by the S2. We note that the critical task of maintaining the precise timing relationship of recorded and reproduced data vs. UTC is the responsibility of the S2 recorder, and requires no special user attention.

3. S2-TCI system design considerations

Due to the high data rate and high data volume implied by the S2 baseband processing system (e.g. a single S2 produces data at a rate of approximately 1 GByte/minute), overall baseband processing efficiency relies on methods of fast data transfer from tape to computer and fast data processing. With current technology, and processing data rates of 128 Mbits/s, it presently appears that even the highest performance Tape-to-Computer Interface and super-computing speeds may barely achieve real-time combined data throughput and processing times. We note that most super-computing applications to-date have been largely compute-intensive rather than data- (and input/output) intensive, but certainly the latter mode is being pursued by the challenging application of pulsar baseband data processing.

Presently, in general the high data rate of the S2 usually implies that data are most efficiently transferred from tape to computer in bursts, and via direct memory access (DMA). In such a scheme, tape data at a given total data rate and for some defined duration are passed to a typically small buffer memory in the DMA system and into computer RAM. Once in RAM, the following processing schemes may be considered:

1. Data are processed immediately from (local) RAM;
2. Data are transferred to an alternative medium, usually disk or output tape device (such as Exabyte or DLT), on the local computer for later processing;
3. Data are transferred to RAM or disk of a more powerful remote computer, such as a super computer, for more efficient processing.

Scheme 1 is generally more appropriate for “smaller-scale” applications requiring “smaller” data volumes and processing times, such as single-pulse studies where data may be processed directly from RAM and only the final results, usually greatly compressed in volume, need to be saved to disk. Even for single-pulse studies, however, reasonable overall processing times demand a capable system both in RAM- and compute-capacity. In this scheme, the entire tape data volume is processed via multiple data transfer bursts to RAM. The primary advantage of this scheme is that the relatively slow transfer of data to disk, typically slower than transfer from S2 to computer RAM, is not needed, and that results are obtained directly. The disadvantage is that generally small data volumes must be processed at one time, and that reasonably large capacity RAM, which is currently cost-prohibitive, is required.

In scheme 2, data are written from RAM to disk or tape before processing. The advantage is that disks or tape devices, currently considerably more cost-effective than RAM, permit a larger data volume to be built up from S2 tape. The disadvantage is that writing to disk followed by subsequent reading from disk for processing generally decreases the overall data processing throughput. Bypassing the disk and writing directly to fast output tape devices (now becoming available) increase the data throughput rate. Furthermore, the advantage of transferring data from S2 tape to more standard tape media/formats such as Exabyte, DAT, DLT, etc., a function which may be performed at established S2 pulsar processing centers, is that S2 data may be made available in industry-standard formats which are generally supported by computers of all kinds. Thus, S2 data processing may effectively be performed at the user’s home institution or (super) computer of choice.

Scheme 3 should be considered when the processing speed in the local computer is limited and a super computer can be used. If a direct interface from S2 tape to super computer is not available, data may be transferred from local RAM or disk across a data link such as ethernet, “fast ethernet”, or the Fiber Distributed Data Interface (FDDI) to the super computer. The advantage of this scheme is the processing speed gain of the super computer. The disadvantage is the requirement for a high-speed link to the super computer. This scheme is clearly feasible only if the processing speed gain in the super computer exceeds the additional link transfer time. The S2-TCI system supports the three schemes discussed above.

In this section, we have described the most general system design considerations, with particular attention to

modes of data transfer from tape recorder to local and remote computers. With regard to the system implementation, we note that this was largely independent of the precise choice of local (workstation) computer, as present technology at the time of writing provided sufficient capabilities (memory sizes, disk transfer speeds, etc.) for the tape-to-computer data transfer. (This implementation was, however, tied to a specific commercial vendor for the DMAC, which at the time of writing necessitated the Sun workstation line of computers.) For example, the first implementation utilized a modest configuration consisting of a “Sparc 20” workstation, with 32 MBytes of RAM, and commercial “disk striping” software to avoid the requirements on special systems (such as “RAID” systems) delivering fast disk speeds. In subsequent implementations regarding the local computer capability, increased RAM and sheer “number crunching” capability were most desired, and primarily for local pulsar signal processing speed in the absence of a more powerful remote computer.

4. S2-TCI software control system

The heart of the S2-TCI system is the control system software, known as *s2tci*. It provides for simple and automated data transfer from S2 tape to computer RAM, and subsequently (nominally) to disk or output tape device. An overview of the *s2tci* control system architecture is shown in Fig. 4. The user specifies transfer parameters such as transfer mode (S2 playback mode), tape start-time and scan length (length of data transfer, e.g. 10 minutes), maximum output file size and output file name(s). *s2tci* parameters may be entered on the command line or read in from an input configuration file. Upon program initiation, *s2tci* checks the user input parameters, available disk space, and other miscellaneous transfer-related issues. It then starts the tapes and checks for valid tape times specified by the user. When a valid transfer configuration is confirmed, *s2tci* begins the data transfer and monitors overall performance. During S2 tape-to-computer data transfer, *s2tci* warns of data transfer difficulties and writes a log file summarizing transfer results. “Bad” portions of tape are flagged and automatically picked up in subsequent “retry passes”. A DMA transfer burst from tape is deemed “bad” and warrants a “retry” if the number of tape (Decoder) sync errors or the estimated percentage of data deemed invalid (both reported by the S2 each second) exceeds thresholds established in the input configuration file.

If data are written to disk or other device, the file format consists of a header containing parameters such as radio source, observing frequency, etc., and data transfer parameters such as S2 tape ID, transfer mode, start time and scan length, followed by the tape (reproduced acquisition system) data bits. A typical individual output file may be 960 MBytes in size, representing one minute of data at a rate of 16 MBytes/s. Various software on- and

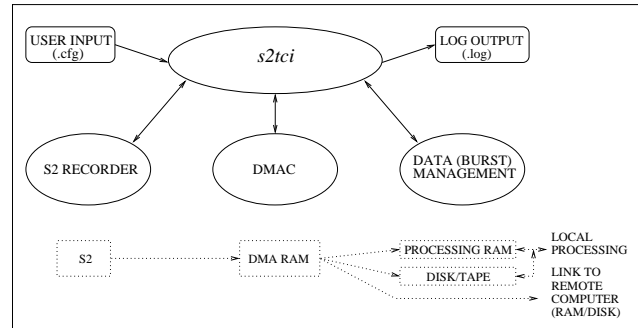


Fig. 4. S2-BPS – Processing Center S2-TCI Control System Overview. In the upper portion of the figure, *s2tci* accepts a user-input transfer configuration file, controls and monitors the S2 recorder and DMA Card (DMAC) systems, manages the DMA data bursts from S2 tape to final destination (“local” or “remote” processing RAM/disk), and writes a log file summarizing the data transfer results. The lower portion of the figure demonstrates the three data management/processing schemes described in Sect. 3. All aspects of S2-TCI operation may be performed as part of an automated unattended supervisory system (even under remote operation), as logically represented in the “Data (Burst) Management” functional block

off-line utilities are provided within *s2tci* for diagnostics and to examine the contents of transferred data files. Data from disk may be written to output tape devices by using standard tape archive programs such as “tar”. They can also be transferred to a more capable computer via fast links such as “fast ethernet” or FDDI if such support exists within the local computer. *s2tci* may be executed from the computer command line, within shell scripts, or within other programs for full flexibility to the user.

5. First results

The first observations using the S2 baseband system for demonstration purposes were made of the Vela pulsar with the 64-m diameter telescope at Parkes, at 6-cm wavelength at left-circular polarization on September 30, 1995. The signal at the telescope was converted to baseband with a 16 MHz bandwidth, sampled at the Nyquist rate with 2-bit quantization, and written to an S2 recorder. At the processing center at ISTS/York University, the data were transferred to a workstation disk via the S2-TCI at the full S2 rate of 16 MBytes/s. The integrated profile of the pulsar is shown in Fig. 5.

6. Conclusions

In this paper, we have concentrated on the technical description and first application of a new type of pulsar instrument for baseband signal acquisition and processing, the S2 baseband processing system (S2-BPS) employing the S2 recording system. By virtue of its low capital and




Fig. 5. The 60 second average pulse profile of the Vela pulsar at 6-cm wavelength shown for a full period. The output of a 16 MHz baseband channel was sampled at the Nyquist resolution of 31.25 ns, de-dispersed, and then smoothed for display purposes to a time resolution of 100 μ s

operations cost and extended recording and playback periods, the S2 has gained widespread acceptance and use in the radio astronomy community. It is currently in use at numerous radio observatories world-wide for VLBI observations. It is also being used for pulsar baseband recording. At the pulsar processing centers, the S2-TCI allows for maximum flexibility in analyzing the pulsar data at the highest theoretical time resolution. A single S2 records 128 Mbits/s and yields 31 ns sample time resolution. Pulsar applications requiring greater bandwidth and higher time resolution can readily use multiple S2's, both during acquisition and processing. Of particular note is the reliability of the S2, which enables tight embedded interaction with the controlling computer, necessary for successful automated, unattended processing of the large data volumes implied by pulsar baseband processing applications. This system has applications in many areas of pulsar research including pulsar timing, pulsar searches, and single pulse studies, and should also find applications in other challenging scientific and instrumentation applications.

Future extensions to the current generation of baseband processing systems, such as the S2-BPS described in this paper, could involve direct interface of tape recorder systems or even radio telescope acquisition systems to super computers, thus eliminate the present bottleneck of the intermediary computer in present implementations. Evolving computer network standards, such as "Gigabit ethernet", will also help in this regard. These systems could implement the next step towards even higher volume and faster processing throughput in the demanding data-intensive applications described here.

Acknowledgements. The authors wish to thank M. Bailes from the Swinburne University of Technology, and J. Reynolds and W. Wilson from the Australian Telescope National Facility for their assistance in acquiring the 1995 Vela pulsar demonstration data from the Parkes radio telescope.

References

- Bartel N., Chandler J.F., Ratner M.I., et al., 1996, AJ 112, 1690
 Cannon W., Baer D., Feil G., et al., 1997, Vistas Astron. 41, No. 2, 297
 Hankins T.H., 1971, ApJ 169, 487
 Hankins T.H., 1996, ASP Conf. Ser. 105 IAU Colloquium 160 Johnson S., Walker M.A., Bailes M. (eds.), p. 197
 Jenet, F.A., Cook, W.R., Prince, T.A., Unwin, S.C., 1997, PASP 109, 707
 Johnson S., Walker M.A., Bailes M., 1996, ASP Conf. Ser. 105 IAU Colloquium 160
 Stinebring D.R., Ryba M.F., Taylor J.H., Romani R.W., 1990, Phys. Rev. Lett. 65, 285
 Taylor J.H., 1991, Proc. I.E.E.E. 79, 1054
 Taylor J.H., 1994, Rev. Mod. Phys. 66, 711
 Wietfeldt R., Van Straten W., Del Rizzo D., et al., 1995, in: Pulsar Science with the Upgraded Arecibo Telescope, 27-28 October, 1995
 Wietfeldt R., Van Straten W., Del Rizzo D., et al., 1996a, ASP Conf. Ser. 105 IAU Colloquium 160 Johnson S., Walker M.A., Bailes M. (eds.), p. 21
 Wietfeldt R.D., Baer D., Cannon W.H., et al., 1996b, IEEE Transactions on Instrumentation and Measurement, Vol. 45, No. 6, 923
 Wolszczan A., 1994, Sci 264, 538