Acceleration Search Methods and Software Improvements for Binary Pulsars (Postprint)

Authors: Ji Changxu, Liu Zhijie, Yu Xuhong

Date: 2021-03-30T00:00:00+00:00

Abstract

Millisecond pulsars are of significant importance for gravitational wave detection, space navigation, and related applications. This paper, by referencing the ATNF database, establishes that a close relationship exists between binary systems and millisecond pulsars. It analyzes the current status of pulsar binary search capabilities at the FAST early science data center. While the Ransom version of the accelsearch algorithm has been highly optimized for CPU, it remains extremely computationally intensive given FAST' s massive data volumes. Lu Jintao' s GPU-based acceleration search achieves a 6-8× speedup, yet fails to effectively optimize GPU memory utilization. This paper first provides an overview of binary pulsars and the effects of their orbital motion on signal processing, then introduces the corresponding search techniques and principles. To enhance the search efficiency for pulsar binaries, we develop programs using parallelization techniques to optimize the GPU version of the acceleration search. Results demonstrate significant speed improvements when processing different .dat files with various zmax parameters.

Full Text

Binary Pulsar System Acceleration Search Method and Software Improvement

Authors: Ji Changxu, Liu Zhijie, Yu Xuhong
Affiliation: Key Laboratory of Information and Computing Science in Guizhou Province

Abstract

Millisecond pulsars hold significant importance for gravitational wave detection and space navigation. By referencing the ATNF database, this paper demonstrates the close relationship between binary star systems and millisecond pulsars. We analyze the current pulse binary search capabilities at the FAST Early Science Data Center. While the Ransom version of the accelsearch algorithm has been highly optimized for CPU, it remains computationally intensive under FAST' s massive data volumes. Although Luo Jintao' s GPU-based acceleration search achieves 6-8× speedup, it fails to optimize GPU memory usage effectively. To improve binary pulsar search efficiency, we first summarize the impact of orbital motion in binary pulsar systems and the corresponding search techniques and principles, then implement parallelization techniques to optimize the GPU version of acceleration search. Results show significant speed improvements across different .dat files using various zmax parameters.

Keywords: Acceleration Search; Binary Pulsar System; FAST

Introduction

The discovery of binary pulsar systems carries profound scientific significance. In 1974, Hulse and Taylor detected a pulsar signal in the Aquila region, which they identified as originating from a binary pulsar (PSR B1913+16). Their in-depth research provided the first indirect quantitative evidence for gravitational waves, earning them the 1993 Nobel Prize in Physics. In February 2020, Wang Lin et al. discovered a millisecond pulsar, M13F, in a binary system within globular cluster M13 using FAST data, while also identifying another pulsar, M13E, as an eclipsing binary. This work marked the first measurement of orbital eccentricities for four binary pulsar systems in M13 and yielded the best timing results for pulsars in M13 to date. These achievements continue to drive scientific progress forward.

Unfortunately, current binary pulsar searches remain time-consuming, and search capabilities urgently require improvement. This challenge manifests in two primary aspects. First, the FAST Early Science Data Center faces demanding requirements for data storage and processing capacity. Following the commencement of FAST' s 19-beam drift-scan survey, data volumes have surged dramatically, generating 50 TB of compressed data daily and reaching 10 PB annually (based on 200 operational days). Effectively utilizing this massive dataset presents a severe test of the data center' s storage and supercomputing capabilities. Second, efficient search capabilities are crucial for FAST to detect high-quality millisecond pulsars. The pulsar distributed search computing system Craber employed at the FAST Early Science Data Center uses the acceleration search method, but search efficiency drops significantly when the zmax parameter is set to large values (typically above 200). Although the Ransom version of accelsearch is highly CPU-optimized, its speed remains inadequate. For instance, processing a single 140 GB drift-scan data file with zmax preset to 200 on an eight-core machine requires approximately 20 hours —and this represents only 10 minutes of surveyed data that has already been processed. Clearly, such data processing capabilities severely limit the effective discovery of binary pulsars.

Moreover, binary systems contain numerous millisecond pulsars. As shown in and [Figure 1: see original paper], globular clusters host substantial populations of binary pulsars, most of which are millisecond pulsars. By the end of 2020, the ATNF pulsar database had cataloged 149 binary pulsars, with approximately 130 discovered in globular clusters. presents sample statistics for several globular clusters, while [Figure 1: see original paper] shows the period distribution for 148 pulsars (one pulsar with a 2.7 s period was excluded as statistically insignificant). Within globular clusters within ~15 kpc of the solar system, binary pulsars constitute a significant fraction—for example, 60% of pulsars discovered in 47 Tuc are binary, and half of Terzan 5' s confirmed pulsars are binaries.

According to the statistical results in [Figure 1: see original paper], millisecond pulsars dominate binary systems. Among the 149 binary pulsars recorded in the ATNF database by December 2020, 113 are millisecond pulsars, accounting for three-quarters (approximately 75.8%) of binary systems. Notably, pulsars with periods of 2–4 ms comprise about 60% of all millisecond pulsars. The close relationship between binary systems and millisecond pulsars was first proposed in the literature, which suggested that millisecond pulsars are normal pulsars spun up to millisecond periods through mass accretion during the low-mass X-ray binary phase. Concurrently, Alpar argued that neutron stars with low magnetic fields accreting matter and spinning up to millisecond periods would evolve into binary systems. Liu et al.' s simulated P–P diagram shows that the spin period distribution of binary pulsar populations peaks around 5 ms, referring to this population as millisecond pulsars. Given this intimate connection and the current time-consuming search methods for binary systems, improving binary pulsar search capabilities is essential. This paper first introduces binary pulsar systems and their search methods, then analyzes these methods and implements parallelization based on their characteristics, achieving several-fold improvements in search efficiency.

1. Binary Pulsar Systems and Acceleration Search

1.1 Binary Systems and Search Techniques

A binary pulsar system consists of two neutron stars orbiting each other, where one is a pulsar and the other is a companion star. Typically, when a star rotates and its magnetic pole beam sweeps across detection equipment, the device receives a pulse signal. Traditional period searches transform time-domain signals to frequency domain through Fast Fourier Transform and identify periodicities by computing spectral power. However, in binary systems, the pulsar' s orbital motion under the companion' s gravitational influence produces acceleration, changing its line-of-sight velocity relative to the detection device. Due to the Doppler effect, the observed spin frequency of the star varies, rendering traditional period search methods ineffective and making pulsar detection difficult.

The acceleration search method effectively eliminates this effect by approximating orbital motion acceleration as constant, thereby removing pulse arrival time variations caused by binary orbital motion in observational data. Under the constant acceleration assumption, the spin frequency drift rate (ḟ) in the observer' s frame relates to acceleration α as:

$$\dot{f} = -\frac{f_0 \alpha}{c}$$

where c is the speed of light and f_0 is the pulsar' s spin frequency in its inertial frame. From a signal processing perspective, methods divide into Time Domain Acceleration Search (TDAS) and Fourier Domain Acceleration Search (FDAS), discussed below.

1.2 Time Domain Acceleration Search

TDAS compensates for Doppler effects caused by orbital motion by linearly stretching or compressing signals to resample data. It iterates through acceleration values within a range to find the time series closest to the true value, then applies period search techniques to identify periodic signals. The time interval for resampling is described by:

$$t'_{sample} = t_{sample}\left(1 + \frac{v}{c}\right)$$

where v is the line-of-sight velocity and c is the speed of light. t_0 is a constant ensuring correct sampling, calculated by:

$$t_0 = \frac{\alpha T^2}{2c}$$

where t_{sample} is the sampling interval, T is the integration time, α is the trial acceleration value, and c is the speed of light. Software packages such as SIGPROC and PEASOUP employ this technique. Using this method, D' Amico et al. discovered four millisecond pulsars in globular clusters, with observational parameters shown in .

1.3 Fourier Domain Acceleration Search

FDAS is based on correlation techniques that enhance spectral power. As mentioned, orbital motion causes observed pulsar signal frequencies to vary, manifesting in the spectrum as power from a particular frequency spreading into adjacent frequency bins, which defeats conventional period search methods. The literature suggests correlating the original signal with neighboring signals to determine similarity and recover the signal, expressed as:

$$A_{r_0} = \sum_{k=-m}^{m} C_k A_{r_0 - k}$$

where r_0 represents the Fourier frequency bin of the pulsar signal being searched (frequency bins in the Fourier domain, represented as 0, 1, 2, 3···for the ith frequency point in the spectrum). A_{r_0} is the Fourier response at frequency bin r_0, which can be recovered through correlation with Fourier responses from m neighboring frequency bins. The asterisk denotes complex conjugation.

This method is generally considered more effective when observation time is much shorter than the orbital period, satisfying:

$$T_{obs} \ll P_{orb}$$

where T_{obs} is the pulsar observation time and P_{orb} is the orbital period. When this condition holds, acceleration can be treated as constant. Combining this with the Doppler formula yields:

$$z = \frac{\alpha f T^2}{c}$$

where c is the speed of light, T is the observation time, f is the harmonic frequency, and z is the number of frequency bins for the assumed frequency drift. Thus, the original signal is recovered through cross-correlation of local signal Fourier responses based on the drift frequency bin count z.

Other period search methods can also address these effects but under different conditions. Phase modulation search primarily targets cases where observation time is comparable to the binary orbital period, while sideband search works better when observation time far exceeds the orbital period. The applicability conditions of acceleration search make it suitable for millisecond pulsars with shorter periods, and in practice, it has discovered many high-quality pulsars. Ransom first used this method to discover pulsar PSR J1807−2459A in globular cluster NGC 6544. Barr E.D. employed this technique to discover the highly eccentric millisecond pulsar PSR J1946+3417, with an eccentricity reaching 0.13-0.14, which is extremely rare and suggests an unusual formation process.

2. Parallelization Necessity and Feasibility Analysis for Fourier Domain Acceleration Search

Parallelization acceleration for FDAS is highly necessary. While some literature argues that orbital motion effects on pulse arrival times are minimal and that binary pulsars can be discovered without acceleration search under certain conditions, the February 2020 discovery of FAST' s first binary pulsar required Dr. Wang Lin et al. to set the zmax parameter to 300 to detect PSR J1641+3627F. Evidently, as FAST' s search requirements continue to escalate, expanding the acceleration search range is essential; otherwise, FAST data cannot be effectively processed, potentially causing us to miss high-quality pulsars.

Compared to TDAS, FDAS is faster and more suitable for parallelization. TDAS requires loading the entire processed time series into memory, whereas FDAS performs convolution on local sequences, making it more appropriate for GPU parallelization with limited memory. When processing different trial acceleration values, TDAS requires FFT transformation for each resampled sequence, while FDAS needs only one FFT transformation followed by correlation operations. Furthermore, FDAS can be viewed as a filtering process, essentially a convolution between the original signal and a filter. Since the original signal length far exceeds the filter length, the computational complexity is approximately $O(N^2)$, where N is the signal length. To improve efficiency, this paper employs Fast Fourier Transform for circular convolution based on a divide-and-conquer approach. The original signal is first divided into segments matching the filter length (with zero-padding in the filter), each segment is convolved with the filter and inverse Fourier transformed to the time domain, and finally the overlapping portions are truncated and merged into the output signal. This process, known as the overlap-save method and illustrated in [Figure 2: see original paper], reduces algorithmic complexity to O(N log N). In summary, the local memory characteristics of this circular convolution-based segmented computation are highly parallelizable and well-suited for parallel computing.

3. Implementation of Fourier Domain Acceleration Search Parallelization

As described in Section 2, the computational bottleneck lies primarily in the numerous parallelizable convolution calculations. Luo Jintao’ s GPU-based acceleration search can be further optimized for memory usage. On one hand, we reduce memory pressure by locally shortening time series lengths through circular convolution. On the other hand, segmented sequences in circular convolution can utilize GPU compute unit local memory for operations. The program flowchart implemented in this paper is shown in Figure 3: see original paper.

When performing Fourier transforms on circularly segmented time series, this paper employs the CUDA library’ s cuFFT interface, which is highly optimized for NVIDIA GPUs. Using batched cuFFT routines, multiple segment Fourier transforms are completed in a single parallel GPU call. For processing acceleration filters, we load the filter within a single GPU Grid, where each thread loads the complete filter, then processes segmented signals post-Fourier transform, performing complex-domain multiplication with different segmented frequency-domain signals. This process occurs locally, with each thread synchronously executing identical operations. The results are then written back to global memory for inverse Fourier transform and overlap removal before merging into the final output.

4. Experiments and Results

Experiments were conducted using an AMD Ryzen7 2700X 8-core 16-thread CPU and two GeForce GTX1080 GPUs. Tests were performed on dedispersed .dat files from FAST drift-scan data (PSR20180811_{00A01H}, PSR20180826_{00A01H}, PSR20180828_{00A01H}, PSR20180902_{00A01H}, with observation parameters in). Different zmax parameters were applied to process these files, generating corresponding ACCEL text files marking periodic signals, with performance compared against both CPU and Luo Jintao' s GPU versions.

The results demonstrate that our parallelized method achieves substantial improvements in acceleration search performance. Compared to the Ransom version, our approach shows an order-of-magnitude speedup of approximately 10-12×, while also outperforming Luo Jintao' s GPU version by about 1.5-1.7×. Overall, the program improvements have proven effective and meet expected performance targets.

References

[1] Hulse R A, Taylor J H. Discovery of a pulsar in a binary system [J]. The Astrophysical Journal, 1975, 195: L51-L53.

[2] J.H.Taylor, L.A.Fowler, P.H.McCulloch, Gao Xuexian, Dai Jingwei. Measuring general relativistic effects in the binary pulsar PSR1913+16 [J]. World Science Translation, 1979(07):21-26.

[3] Wang L, Peng B, Stappers B W, et al. Discovery and Timing of Pulsars in the Globular Cluster M13 with FAST[J]. The Astrophysical Journal, 2020, 892(1).

[4] Zhang H, Xie X, Li D, et al. A PB-scale pulsar data processing acceleration method and system for FAST[J]. Astronomical Research & Technology, 2020, 18(1): 129-137.

[5] Pan Z., Hobbs G., Li D., Ridolfi A., Wang P., Freire P.. Discovery of two new pulsars in 47 Tucanae (NGC 104)[J]. Monthly Notices of the Royal Astronomical Society: Letters, 2016, 459(1).

[6] Zhichen Pan, Scott M. Ransom, Duncan R. Lorimer, William C. Fiore, Lei Qian, Lin Wang, Benjamin W. Stappers, George Hobbs, Weiwei Zhu, Youling Yue, Pei Wang, Jiguang Lu, Kuo Liu, Bo Peng, Lei Zhang, Di Li. The FAST Discovery of an Eclipsing Binary Millisecond Pulsar in the Globular Cluster M92 (NGC 6341)[J]. The Astrophysical Journal Letters, 2020, 892(1).

[7] Andersen B C, Ransom S M. A Fourier Domain "Jerk" Search for Binary Pulsars[J]. The Astrophysical Journal Letters, 2018, 863(1): L13.

[8] Ransom S M, Greenhill L J, Herrnstein J R, et al. A binary millisecond pulsar in globular cluster NGC 6544[J]. The Astrophysical Journal Letters, 2000,

546(1): L2.

[9] Lorimer D R. Binary and millisecond pulsars at the new millennium[J]. Living reviews in relativity, 2001, 4(1): 1-87.

[10] Backer D C, Kulkarni S R, Heiles C, et al. A millisecond pulsar[J]. Nature, 1982, 300(5893): 615-618.

[11] Alpar M A, Cheng A F, Ruderman M A, et al. A new class of radio pulsars[J]. Nature, 1982, 300(5894): 728-730.

[12] Liu Peng, Wang Pei, Li Di, Zhang Jie, Zhang Lei, Zhang Chengmin, Zhu Weiwei, Yue Youling, Dai Shi. FAST 19-beam pulsar drift-scan survey simulation[J]. Progress in Astronomy, 2018, 36(02):173-188.

[13] Li Xiangdong, Wang Zhenru. Observational characteristics of X-ray binary pulsars[J]. Progress in Astronomy, 1993(02):106-116.

[14] Shuai Ping, Chen Shaolong, Wu Yifan, et al. Principle of X-ray pulsar navigation[J]. Journal of Astronautics, 2007, 28(6): 1538.

[15] Pan Zhichen, Qian Lei, Yue Youling. Pulsar search techniques and prospects for FAST pulsar searches[J]. Astronomical Research & Technology, 2017, 14(01):8-16.

[16] Cameron A D. Innovative Pulsar Searching Techniques[J]. Williams College, 2018.

[17] Lorimer D R. SIGPROC: pulsar signal processing programs[J]. Astrophysics Source Code Library, 2011, 1107.016.

[18] D' Amico N, Lyne A G, Manchester R N, et al. Discovery of short-period binary millisecond pulsars in four globular clusters[J]. The Astrophysical Journal Letters, 2001, 548(2): L171.

[19] Ransom S M, Eikenberry S S, Middleditch J. Fourier techniques for very long astrophysical time-series analysis[J]. The Astronomical Journal, 2002, 124(3): 1788.

Note: Figure translations are in progress. See original paper for figures.