

UWLPIPE: Ultra-wide Bandwidth Low-frequency Pulsar Data Processing Pipeline (Postprint)

Authors: Ya-Zhou Zhang, Hai-Long Zhang, Jie Wang, Jian Li, Xin-Chen Ye, Shuang-Qiang Wang, Xu Du, Han Wu, Ting Zhang and Shao-Cong Guo

Date: 2024-08-14T00:00:00+00:00

Abstract

For real-time processing of ultra-wide bandwidth low-frequency pulsar baseband data, we designed and implemented an ultra-wide bandwidth low-frequency pulsar data processing pipeline (UWLPIPE) based on the shared ringbuffer and GPU parallel technology. UWLPIPE runs on the GPU cluster and can simultaneously receive multiple 128 MHz dual-polarization VDIF data packets preprocessed by the front-end FPGA. After aligning the dual-polarization data, multiple 128M subband data are packaged into PSRDADA baseband data or multi-channel coherent dispersion filterbank data, and multiple subband filterbank data can be spliced into wideband data after time alignment. We used the Nanshan 26 m radio telescope with the L-band receiver at 964–1732 MHz to observe multiple pulsars. Finally, we processed the data using DSPSR software, and the results showed that each subband could correctly fold out the pulse profile, and the wideband pulse profile accumulated by multiple subbands could be correctly aligned.

Full Text

Preamble

ChinaXiv Research in Astronomy and Astrophysics, 24:075011 (13pp), 2024 July

© 2024. National Astronomical Observatories, CAS and IOP Publishing Ltd. Printed in China and the U.K.

<https://doi.org/10.1088/1674-4527/ad4fc4>

UWLPIPE: Ultra-wide Bandwidth Low-frequency Pulsar Data Processing Pipeline

Ya-Zhou Zhang^{1,2}, Hai-Long Zhang^{1,2,3,4}, Jie Wang^{1,4}, Jian Li^{1,2}, Xin-Chen

Ye^{1, 2, 4}, Shuang-Qiang Wang¹, Xu Du^{1, 2}, Han Wu^{1, 2}, Ting Zhang^{1, 2}, and Shao-Cong Guo⁵

¹ Xinjiang Astronomical Observatory, Chinese Academy of Sciences, Urumqi 830011, China; zhanghailong@xao.ac.cn

² University of Chinese Academy of Sciences, Beijing 100049, China

³ Key Laboratory of Radio Astronomy, Chinese Academy of Sciences, Nanjing 210008, China

⁴ Xinjiang Key Laboratory of Radio Telescope Technology, Urumqi 830011, China

⁵ Southeast University, Nanjing 211189, China

Received 2024 April 16; revised 2024 May 14; accepted 2024 May 21; published 2024 July 8

Abstract

To address the real-time processing requirements for ultra-wide bandwidth low-frequency pulsar baseband data, we designed and implemented an ultra-wide bandwidth low-frequency pulsar data processing pipeline (UWLPIPE) based on shared ringbuffer and GPU parallel technology. UWLPIPE runs on a GPU cluster and can simultaneously receive multiple 128 MHz dual-polarization VDIF data packets preprocessed by the front-end FPGA. After aligning the dual-polarization data, multiple 128 MHz subband data streams are packaged into either PSRDADA baseband data or multi-channel coherent dedispersion filterbank data, with multiple subband filterbank data streams subsequently spliced into wideband data after time alignment. We used the Nanshan 26 m radio telescope with the L-band receiver operating at 964–1732 MHz to observe multiple pulsars. The data were processed using DSPSR software, and the results demonstrated that each subband could correctly produce a folded pulse profile, with the wideband pulse profile accumulated from multiple subbands achieving proper alignment.

Key words: (stars:) pulsars: general – methods: data analysis – techniques: miscellaneous

1. Introduction

With the rapid development of key technologies, the performance requirements for digital backend systems in astronomical observations continue to increase. Backend systems must perform high-speed sampling, real-time analysis, and data preprocessing under wider bandwidths, higher time resolution, and higher frequency resolution. The deployment of Phased Array Feeds (PAF) (van Cappellen et al. 2022; Zhang & Duan 2022) and multi-beam receiving systems (Yang & Han 2022) has doubled the information obtained from astronomical observations. The high-speed data streams generated during transmission, storage, and real-time processing on heterogeneous platforms present urgent challenges

for the operation of various radio observation instruments. Wider sampling bandwidths, higher digital signal bit widths, and more array antennas lead to exponential growth in data volume. Due to storage device performance limitations, massive astronomical signals can only be processed and analyzed in real time, posing significant challenges to computing hardware and software (Wei 2019).

To address high-speed data transmission and processing, a viable solution involves creating shared ringbuffers in memory for real-time data caching and transferring data to GPUs for immediate CPU+GPU processing. Internationally, heterogeneous computing platforms for real-time processing of pulsar signals have become mainstream (Wei et al. 2023). Utilizing GPU computational units to process data can reduce CPU load and significantly enhance system data stream processing efficiency. With continuous development of digital backend technology in radio astronomy, real-time signal processing poses severe challenges to traditional computing technologies, while the powerful computational capabilities of GPU clusters offer a feasible solution for handling massive data streams generated during radio astronomical observations.

Currently, mainstream digital backend systems often adopt a hybrid architecture of Field Programmable Gate Array (FPGA), CPU, and GPU. Heterogeneous systems require data exchange between different devices. After preprocessing by FPGA, data is transmitted via network to the CPU memory of the data processing server and then copied to GPU memory for further processing. One key technical challenge for data processing systems based on heterogeneous platforms is achieving high-speed data circulation between FPGA and CPU, and between CPU and GPU.

In radio astronomy digital backend systems, the program that achieves high-rate real-time transmission and preprocessing of data packets is called a pipeline (Paine & Lee 2014), which includes several steps such as reception, data buffering, data preprocessing, and packaging into astronomical standard formats.

[Figure 1: see original paper]. UWLPIPE data flow.

[Figure 2: see original paper]. Ringbuffer cyclic read and write operations.

To address quasi-real-time or real-time high-speed transmission and preprocessing of radio astronomical data, several pipeline software packages have been developed both domestically and internationally. GUPPI_{daq} (Data acquisition software for GUPPI) is the data acquisition software for the Green Bank Telescope pulsar digital backend Green Bank Ultimate Pulsar Processing Instrument (GUPPI). It operates in multi-thread mode, with separate threads for network data reception and data processing that communicate through shared memory ringbuffers. After preprocessing real-time network data streams, GUPPI_{daq} stores them in PSRFITS format (Hobbs 2021) in either fold or search mode based on configuration.

High Availability SHared PIPEline Engine (HASHPIPE) is an efficient data processing framework based on shared memory. The core of HASHPIPE is

the shared ringbuffer, with applications written as shared library “plugins” that enable data circulation and sharing among multiple threads. The shared memory buffer between threads is controlled by semaphores to manage tasks, achieving mutual exclusion between read and write threads and preventing data read/write errors (MacMahon et al. 2018).

Niu et al. (2019) designed the Tianlai Disk Array Correlator based on Reconfigurable Open Architecture Computing Hardware-2 (ROACH2) and GPU, implementing UDP network data reception and subsequent data processing in GPU servers. Pei et al. (2022) developed HRBF_{HASHPIPE} based on HASHPIPE for PAF systems. After FPGA collected simulated signals, data with 256 MHz bandwidth was output to a GPU server through four 10 Gigabit Ethernet (10GbE) links. The server ran four instances of HRBF_{HASHPIPE}, each receiving 64 MHz bandwidth data and invoking a GPU card for beamforming calculations. According to the future plan for the QiTai radio Telescope (QTT; Wang et al. 2023), Pei et al. (2023) designed UWB_{HASHPIPE} to multi-threadedly receive and store multiple subband data in parallel, which could be flexibly configured for data distribution links based on IP addresses and port numbers. UWB_{HASHPIPE} was tested in the Nanshan 26 m radio telescope (NSRT) pulsar observation experiment, where the collected data bandwidth was 512 MHz transmitted to two servers via eight links. Each server ran four instances of UWB_{HASHPIPE} for data processing and signal encapsulation. The test showed that the acquisition system had high precision, good data integrity, and the signal-to-noise ratio of the merged pulsar profile was superior to that of single subband data.

PSRDADA (Distributed Acquisition and Data Analysis for Radio Astronomy) is capable of flexibly managing a shared ringbuffer, which is used for distributed recording and processing of pulsar baseband data (van Straten & Jameson 2021). Based on PSRDADA, it is possible to implement read-write clients that write data into the ringbuffer and read data from it. The configuration and control of coherent processes are initiated through a web-based interface that launches scripts. Data recording backends such as the ATNF Parkes Swinburne Recorder (APSR), Berkeley Parkes Swinburne Recorder (BPSR), and CASPER Parkes Swinburne Recorder (CASPSR) at the Parkes radio telescope are all developed based on PSRDADA.

The Parkes ultra-wide bandwidth low-frequency (UWL) receiving system employs an FPGA+GPU architecture for preprocessing and recording observational data. It divides the 704–4032 MHz wideband signal into three analog RF bandwidths for sampling. During the FPGA preprocessing stage, the Polyphase Filter Bank (PFB; Harris & Haines 2011) channelization technique is used to further divide it into 26 continuous subbands, each with a bandwidth of 128 MHz. The FPGA packages the data in VLBI Data Interchange Format (VDIF; Whitney et al. 2010) and transmits it to various GPU nodes via UDP protocol. To cope with high-speed network streams, a high-speed ringbuffer is established using PSRDADA on each GPU node to temporarily store data. Subsequently,

the data is copied from the buffer to GPU memory for further processing, which can include operations such as pulsar folding, pulsar searching, Fast Radio Burst searching, and spectral line observations (Hobbs et al. 2020).

[Figure 3: see original paper]. Observation configuration file.

[Figure 4: see original paper]. Machine configuration file.

2. UWLPIPE

Ultra-wide bandwidth low-frequency pulsar data processing pipeline (UWLPIPE) is a simple and efficient pulsar backend processing pipeline software. It can quickly configure the current observation by reading configuration information from a TOML file. For example, it can obtain observation source information (pulsar source name, DM, observation bandwidth, observation center frequency, etc.) and configure the GPU server network (IP and port, etc.). As shown in Figure 1, by receiving dual-polarization VDIF packets and unpacking them according to the VDIF header information, the data is rearranged based on the polarization timestamp information. Afterwards, the “write” client is enabled to write the data into the shared ringbuffer. The “read” client reads the data from the buffer, copies it to GPU memory for processing, and finally encapsulates it into standard astronomical data format.

Shared memory can serve as a communication medium between processes for synchronous data exchange. The write client process writes data into shared memory, while the read client reads data from the shared buffer, processing it differently according to various research needs. As shown in Figure 2, the ringbuffer is essentially linear memory that is logically connected at the head and tail, facilitating cyclic read and write operations on the buffer.

From the perspective of buffer operations, processes can be roughly divided into two modes: reading and writing. In writing mode, the client receives high-speed network data and writes it into the shared buffer through rearrangement. There might be various types of reading clients, such as the baseband data encapsulation process that reads data from the shared buffer, adds header information, and directly stores it on disk. Another example is the filterbank channelization process, which reads data from the shared buffer and copies it to GPU memory for subsequent pipeline processing, such as channelization, coherent dedispersion, integration, and folding. Other operations include RFI mitigation, pulsar searching, and VLBI. The pulsar data processing pipeline based on shared memory allows for modular code writing according to specific functions, ensuring good extensibility.

2.1. UWLPIPE Configuration

Currently, UWLPIPE configures observation targets and GPU nodes by reading the configuration files `Observation.toml` and `Machine.toml`. As shown in Figure 3, `Observation.toml` includes information such as the pulsar name and disper-

sion measure, telescope name, receiver name, number of observation subbands, bandwidth size, center frequency of each subband, and storage format.

Machine.toml is shown in Figure 4, which allows configuration of IP and port for receiving subband network data, as well as the key values, number, and size of the circular buffer. Since baseband data storage can be selected, there are certain requirements for the speed of the storage medium. The data of each subband is stored on a separate high-speed solid-state disk.

2.2. High-speed Network Data Reception and Unpacking Algorithm

Traditional TCP/IP technology, in the process of handling massive astronomical data packets, must go through the operating system and other software layers, which requires substantial server resources and memory bus bandwidth. Data is copied and moved back and forth between system memory, processor cache, and network controller cache, imposing a heavy burden on the server's CPU and memory and exacerbating network latency effects.

To reduce system overhead brought about by the transmission of massive astronomical data, we reconstructed the high-speed network unpacking program using the VMA library. The VMA library utilizes RDMA-enabled network cards for direct hardware access and advanced polling techniques to achieve kernel bypassing. This allows the VMA library to bypass the kernel's network stack for all IP network traffic transmission and reception socket API calls, reducing CPU usage, alleviating memory bandwidth bottlenecks, and enhancing bandwidth utilization efficiency.

The future planning of QTT's UWL receiving system is the same as Australia's Parkes, with a bandwidth range of 704–4032 MHz (Wang et al. 2023). During the FPGA preprocessing stage, PFB technology is used to divide the 704–4032 MHz bandwidth signal, totaling 3328 MHz, into 26 dual-polarization subband signals of 128 MHz bandwidth each. The data quantization precision is 32 bit (real part 16 bit + imaginary part 16 bit). As shown in Figure 5, the VDIF header information is depicted, with key and fixed information highlighted in red. Other channel numbers are set to 1, with $\log_2(\text{nchan}) = 0$. Each VDIF data frame size is 8224 bytes (header 32 bytes + data 8192 bytes), set to $8224/8 = 1028$. The complex flag bit is set to 1, with a quantization precision of 32 bit, set to $\text{nbit} = 31$. Dual-polarization data is distinguished by Thread ID, where 0 represents polarization 0 and 1 represents polarization 1. The number of data frames per second is 62,500 (0–62,499), calculated according to Equation (1), where f represents the sampling rate of 128 MHz and nbit represents the quantization precision of 32.

$$n_{\text{frame}} = \frac{f_s \times 2}{8192/(\text{nbit}/8)} = 62,500$$

After packaging each channel of dual-polarization subband data into VDIF for-

mat, FPGA transmits it to the GPU server via a high-speed network. Each channel of data is sent to the corresponding 100GbE network card on the GPU, and the GPU server receives data by listening to specific IP addresses and ports.

Since the UDP transmission protocol is adopted, processing is required for lost, disordered, and duplicate packets; otherwise, phase information will be missing. For pulsar data, packet loss can severely affect pulsar period prediction and folding. Astronomical observation has high requirements for time accuracy, so it is necessary to align dual-polarization data. To solve these problems, we propose a Linked-list Double-buffer Receive Packet (LDRP) algorithm.

LDRP is a dual-polarization VDIF packet alignment algorithm based on a double buffer structure for alternate storage. As shown in Figure 6, two consecutive memory segments are allocated as buffers, named curr_buffer and next_buffer . Three byte offset pointers are set up as curr_byte_offset , next_byte_offset , and last_byte_offset . Their initial values are set as shown in Equation (2), pointing to the start of curr_buffer , the end of curr_buffer (the start of next_buffer), and the end of next_buffer , respectively.

$$\begin{aligned}\text{curr_byte_offset} &= \text{curr_byte_offset} \\ \text{next_byte_offset} &= \text{curr_byte_offset} + \text{bufsize} \\ \text{last_byte_offset} &= \text{next_byte_offset} + \text{bufsize}\end{aligned}$$

Before performing polarization alignment, it is necessary to set a reference start time. The header of the VDIF data frame contains an accurate timestamp and counter, which can be converted into the byte offset from the reference start time. When starting to receive network data packets, unpack the first VDIF packet received for polarization 0 to obtain the timestamp seconds of the data packet. If the packet's counter is greater than 0, add 1 to the timestamp seconds as the reference time, with the aim of starting to receive data at the beginning of an entire second. Subsequently received VDIF packets can calculate the byte offset from the reference start time using Equation (3), as shown in Figure 7.

$$\text{byte_offset} = \text{offset} \times \text{bytes} + \frac{\text{iframe}}{\text{frames}} \times \text{bytes} + \text{thread_id} \times \frac{\text{bytes}}{2}$$

where offset equals the difference between the VDIF packet's timestamp seconds and the reference start time, bytes represents the number of bytes sent per second, iframe is the counter number of the VDIF packet, frames is twice the number of data bytes in the VDIF packet, thread_id equals 0 for polarization 0 and 1 for polarization 1, and bytes is the size of the VDIF packet in bytes. The LDRP algorithm flow is shown in Figure 8. The detailed processing of VDIF packets is shown in Figure 9, which handles several cases based on byte_offset .

1. When the offset is less than 0, the timestamp of the VDIF data packet is before the reference start time, so it is discarded directly, as shown in the dotted white box in Figure 9.
2. When byte_offset is greater than or equal to $\text{curr_}\{\{\text{byte}\}\}\{\{\text{offset}\}\}$ and less than $\text{next}\{\{\{\text{byte}\}\}\}\{\{\text{offset}\}\}$, it indicates that the VDIF data packet belongs to the $\text{curr}\{\text{buffer}\}$ buffer. The difference between byte_offset and $\text{curr_}\{\{\text{byte}\}\}\{\{\text{offset}\}\}$ is calculated as the offset position of the VDIF data block relative to the starting position of $\text{curr}\{\text{buffer}\}$, and the data is copied to the corresponding position in $\text{curr_}\{\text{buffer}\}$, with the counter of $\text{curr_}\{\text{buffer}\}$ incremented by 1.
3. When byte_offset is greater than or equal to $\text{next_}\{\{\{\text{byte}\}\}\}\{\{\text{offset}\}\}$ and less than $\text{last}\{\{\{\text{byte}\}\}\}\{\{\text{offset}\}\}$, it indicates that the VDIF data packet belongs to the $\text{next}\{\text{buffer}\}$ buffer. The difference between byte_offset and $\text{next_}\{\{\{\text{byte}\}\}\}\{\{\text{offset}\}\}$ is calculated as the offset position of the VDIF data block relative to the starting position of $\text{next}\{\text{buffer}\}$, and the data is copied to the corresponding position in $\text{next_}\{\text{buffer}\}$, with the counter of $\text{next_}\{\text{buffer}\}$ incremented by 1.
4. When byte_offset is greater than or equal to $\text{last_}\{\{\{\text{byte}\}\}\}\{\{\text{offset}\}\}$, it indicates an abnormality in the network data (high system load or other network issues), and the VDIF packet has exceeded the range of the double buffer. In response to this situation, as shown in Figure 10, it is necessary to immediately update the pointers of the double buffer structure to avoid further packet loss.
5. When the $\text{curr_}\{\text{buffer}\}$ buffer is filled with data (the counter of $\text{curr_}\{\text{buffer}\}$ equals the buffer capacity), it indicates that the data in the $\text{curr_}\{\text{buffer}\}$ buffer is complete and can be processed by subsequent steps. At this time, it is necessary to update the double buffer offset pointers and swap buffers, as shown in Figure 10.
6. When the counter of $\text{next_}\{\text{buffer}\}$ is greater than or equal to half the capacity of $\text{next_}\{\text{buffer}\}$, it indicates that the data in the $\text{curr_}\{\text{buffer}\}$ buffer is not filled, suggesting VDIF data packet loss. To avoid further packet loss, it is necessary to update the double buffer offset pointers and swap buffers, as shown in Figure 10.

For ultra-wide bandwidth pulsar baseband data, tens of thousands of VDIF data packets are received per second, and the time interval t between VDIF data packets is very small, as shown in Figure 11. When one of the double buffers is filled with data, it needs to be copied to a shared buffer or directly to GPU memory for subsequent processing. If the copy time is greater than the time interval t , it will cause loss of subsequent VDIF data packets due to the inability to receive them in time. To address this issue, we employ multi-threaded asynchronous copying. Specifically, when the buffer is filled with data, a new thread is initiated to copy the data to the location required for subsequent

steps, while the main receiving thread continues to receive VDIF data packets. Both processes proceed in parallel without interfering with each other.

2.3. Pulsar Data Processing and Astronomical Data Format Packaging Technology

During the FPGA preprocessing stage, the wideband signal is divided into multiple narrowband signals using PFB/Oversampled Polyphase Filter Banks (OPFB; Zhang et al. 2023b), which are then transmitted to various nodes of the GPU cluster for processing. UWLPIPE is capable of handling subband data from both PFB and OPFB channel division modes. For OPFB (with a 4/3 oversampling factor), the data bandwidth size for each subband transmission is 128 MHz. After performing FFT operations, it is necessary to discard data corresponding to the first and last 1/6 of the bandwidth, retaining the useful data within the central 128 MHz bandwidth. Subsequent steps only process this 128 MHz data, ultimately encapsulating it in an astronomical data format. However, if encapsulated as DADA baseband data, to reduce computational time, the data is stored directly at 128 MHz to minimize processing time.

2.3.1. Baseband Data Record UWLPIPE, based on configuration information, is capable of directly packaging data into DADA baseband data format. This DADA baseband data is identical to the baseband data generated by the Parkes Medusa backend and can be processed by DSPSR (van Straten & Bailes 2011). For PFB/OPFB sub-channel technology, the key header information for the generated DADA baseband is shown in Table 1.

Key header information for DADA baseband data generated by PFB/OPFB

Keyword	Value
BW	Bandwidth size
CFREQ	Center frequency
HDR_{SIZE}	Header size
HDR_{VERSION}	Header version
NBIT	Quantization accuracy of recorded data
NCHAN	Number of channels
NDIM	1 for real numbers, 2 for complex numbers
NPOL	Number of polarizations
RECEIVER	Receiver type (set to UWL)
TSAMP	Sampling interval time (s)
UTC_{START}	Start time of baseband record (yyyy-MM-dd-HH:mm:ss)

Among these, BW represents the bandwidth size, CFREQ represents the center frequency, NBIT represents the quantization accuracy of the recorded data, NDIM is 1 for real numbers and 2 for complex numbers, NPOL represents the number of polarizations, RECEIVER represents the type of receiver (here set

to UWL, which DSPSR can recognize as baseband data recorded by ultra-wide bandwidth receiver backend), TSAMP represents the interval time of sampling points, and UTC_{START} is the starting time of the baseband record.

The data portion needs to be encoded using offset binary, as shown in Equation (4):

$$\text{output} = \text{input} \oplus 0x8000$$

where \oplus represents the XOR operation symbol, input represents the input data, and output represents the encoded data. For dual-polarization data, since the data portion in the VDIF data frame only contains one polarization's data, and the size is 8192 bytes, under 16-bit quantization precision, the two polarizations in the baseband data are alternately stored every 2048 numbers.

2.3.2. Filterbank Data Record UWLPIPE can package data into filterbank format according to configuration. For the filterbank format, it is necessary to further divide each subband into channels—into 128 subbands (each subband has a bandwidth of 1 MHz)—and perform coherent dedispersion on each subband. As shown in Figure 12, after performing FFT on the time data sequence, it is further channelized into multiple subbands. According to the center frequency, bandwidth, and dispersion measure (DM) of each subband, the chirp coefficient (the inverse function of the interstellar medium transfer function) is generated (Zhang et al. 2023a), as shown in Equation (5):

$$\text{chirp} = e^{-2\pi i \cdot \text{DM} \cdot \frac{f_c^2}{f^2}}$$

After coherent dedispersion, in order to reduce the data volume, the data is integrated (to 32 s) and finally stored as filterbank data (Lorimer 2011).

UWLPIPE will eventually output multiple filterbank subbands with a bandwidth of 128 MHz. Due to network latency differences among GPU nodes, the starting record times of each subband may vary. In order to synthesize a wideband signal, it is necessary to align the times of multiple subbands. As shown in Figure 13, find the maximum start time among the multiple subbands, use it as the effective reference start time for synthesizing the wideband, skip a certain number of bytes in other subbands to align with the reference time, and synthesize the final wideband signal.

3. Results

UWLPIPE was tested using the L-band 1 GHz bandwidth receiver of NSRT. The test scheme is shown in Figure 14, which used three FPGA development boards. The backend uses two GPU servers for data processing; Table 2 shows the configuration of each GPU server. Each development board outputs two

128 MHz bandwidth dual-polarization subband data streams, with a total bandwidth range of 964–1732 MHz, covering 768 MHz.

GPU Server Configuration

Component	Specification
Operating System	Debian 10
CPU	2× Intel Xeon Gold 5317
GPU	3× NVIDIA GeForce RTX 3090
Memory Size	512 GiB
Storage Size	3× NVME SSD (8 TB)

Based on the PFB subband technique, UWLPIPE observed J0332+5434 for 5 minutes. After processing with DSPSR, the dynamic spectra of the six subbands are shown in Figure 15. The final synthesized wideband signal's dynamic spectrum after processing is shown in Figure 16. Each subband correctly produces a folded pulse profile, and the final synthesized wideband signal's processed pulse profile aligns accurately. The frequency spectrum of the six subbands is shown in Figure 19(a), which clearly shows significant attenuation between subbands.

Using the OPFB (Orthogonal Polyphase Filter Banks) subband technique, UWLPIPE also observed J0332+5434 for 5 minutes. The dynamic spectra of the six subbands are shown in Figure 17, and the final synthesized wideband signal's dynamic spectrum after processing is shown in Figure 18. The frequency spectrum of the six subbands is shown in Figure 19(b), and when compared to Figure 19(a), it is evident that there is no attenuation between OPFB subbands, which does not affect the continuity of the pulsar signal.

Furthermore, UWLPIPE was used to observe J1935+1616 and J2022+5154. The dynamic spectra of the processed six subbands are shown in Figures 20 and 21, respectively. The final synthesized wideband signals' dynamic spectra after processing are shown in Figures 22 and 23, respectively. The results demonstrate that UWLPIPE can correctly receive subband data from both PFB and OPFB channels, verifying the effectiveness of UWLPIPE in data processing.

[Figure 14: see original paper]. NSRT test architecture.

[Figure 15: see original paper]. J0332+5434's subband Dynamic Spectrum Diagram (PFB).

[Figure 16: see original paper]. Dynamic spectrum of J0332+5434 (PFB).

[Figure 17: see original paper]. J0332+5434's subband Dynamic Spectrum Diagram (OPFB).

[Figure 18: see original paper]. Dynamic spectrum of J0332+5434 (OPFB).

[Figure 19: see original paper]. J0332+5434 frequency spectrum.

[Figure 20: see original paper]. J1935+1616's subband Dynamic Spectrum Diagram (OPFB).

[Figure 21: see original paper]. J2022+5154's subband Dynamic Spectrum Diagram (OPFB).

[Figure 22: see original paper]. Dynamic spectrum of J1935+1616 (OPFB).
[Figure 23: see original paper]. Dynamic spectrum of J2022+5154 (OPFB).

4. Summary

To address the real-time processing requirements for ultra-wide bandwidth pulsar data, we designed and implemented UWLPIPE based on GPU parallel technology. UWLPIPE can configure the current observation source and GPU node parameters by reading TOML configuration file information. UWLPIPE demonstrates proficiency in real-time parallel reception of multi-subband dual-polarization pulsar signals. It facilitates multi-channel parallel real-time coherent dedispersion processing, enables ultra-wide bandwidth data synthesis, and packages data into formats suitable for scientific data processing needs.

The LDRP algorithm is designed and implemented to unpack the header information of dual-polarization VDIF data packets, extracting timestamp information. By calculating the offset of each VDIF data packet relative to the reference time, the data is placed in the correct position within a double buffer structure to align the dual-polarization data.

UWLPIPE is capable of packaging data into PSRDADA baseband data and filterbank data astronomical formats—that is, directly saving the aligned dual-polarization data as baseband data, or storing filterbank data after performing operations such as channel separation, coherent dedispersion, and integration on the GPU.

We tested UWLPIPE using NSRT with an L-band bandwidth range of 964–1732 MHz. During the FPGA preprocessing stage, we used PFB and OPFB channel separation respectively for 5-minute observations of J0332+5434. Comparing the spectra across six subbands, it is evident that the OPFB-based channel separation technique yields a smoother spectrum compared to PFB, eliminating the spectral decay effect between subbands.

We also observed J1935+1616 and J2022+5154, obtaining six subbands of 128 MHz data. After performing folding integration on each subband, we obtained correct pulse profiles. Finally, we merged the six subbands to generate wideband data and performed folding integration again. The generated subband profiles were correctly aligned. These experiments verified the accuracy of UWLPIPE's data processing and laid the foundation for real-time processing of ultra-wide bandwidth pulsar data.

Acknowledgments

This work is supported by the National Key R&D Program of China Nos. 2021YFC2203502 and 2022YFF0711502; the National Natural Science Foundation of China (NSFC) (12173077); the Tianshan Talent Project of Xinjiang Uygur Autonomous Region (2022TSYCCX0095 and 2023TSYCCX0112); the Scientific Instrument Developing Project of the Chinese Academy of Sciences,

grant No. PTYQ2022YZZD01; China National Astronomical Data Center (NADC); the Operation, Maintenance and Upgrading Fund for Astronomical Telescopes and Facility Instruments, budgeted from the Ministry of Finance of China (MOF) and administrated by the Chinese Academy of Sciences (CAS); and the Natural Science Foundation of Xinjiang Uygur Autonomous Region (2022D01A360).

ORCID iDs

Ya-Zhou Zhang: <https://orcid.org/0000-0001-6046-2950>

Hai-Long Zhang: <https://orcid.org/0000-0002-8951-7094>

Jie Wang: <https://orcid.org/0000-0003-0380-6395>

Xu Du: <https://orcid.org/0000-0001-6448-0822>

References

- DuPlain, R., Ransom, S., & Demorest, P. 2008, *Proc. SPIE*, 7019, 70191A
- Harris, C., & Haines, K. 2011, *PASA*, 28, 317
- Hobbs, G. 2021, *pfits: PSRFITS-format Data File Processor*, *Astrophysics Source Code Library*, ascl:2104.013
- Hobbs, G., Manchester, R. N., Dunning, A., et al. 2020, *PASA*, 37, e012
- Lorimer, D. R. 2011, *SIGPROC: Pulsar Signal Processing Programs*, *Astrophysics Source Code Library*, ascl:1107.016
- MacMahon, D. H. E., Price, D. C., Lebofsky, M., et al. 2018, *PASP*, 130, 5002
- Niu, C.-H., Wang, Q.-X., MacMahon, D., et al. 2019, *RAA*, 19, 102
- Paine, D., & Lee, C. P. 2014, in *2014 IEEE 10th International Conf. e-Science*, Vol. 1 (Sao Paulo: IEEE), 231
- Pei, X., Li, J., Duan, X., & Zhang, H. 2023, *PASP*, 135, 075003
- Pei, X., Wang, N., Werthimer, D., et al. 2022, *RAA*, 22, 045016
- van Cappellen, W. A., Oosterloo, T. A., Verheijen, M. A. W., et al. 2022, *A&A*, 658, A146
- van Straten, W., & Bailes, M. 2011, *PASA*, 28, 1
- van Straten, W., Jameson, A., & Osłowski, S. 2021, *PSRDADA: Distributed Acquisition and Data Analysis for Radio Astronomy*, *Astrophysics Source Code Library*, ascl:2110.003
- Wang, N., Xu, Q., Ma, J., et al. 2023, *SCPMA*, 66, 289512
- Wei, D. 2019, *Research on Key Technologies in Radio Astronomical Data Real-time Processing*, PhD thesis, University of Chinese Academy of Sciences
- Wei, J., Zhang, C., Zhang, Z., et al. 2023, *SSPMA*, 53, 229506
- Whitney, A., Kettenis, M., Phillips, C., & Sekido, M. 2010, in *Sixth Int. VLBI Service for Geodesy and Astronomy. Proc. 2010 General Meeting*, ed. R. Navarro et al., 192
- Yang, J., & Han, W.-l. 2022, *ChA&A*, 46, 309
- Zhang, H.-L., Zhang, Y.-Z., Zhang, M., et al. 2023a, *RAA*, 23, 015023
- Zhang, M., Zhang, H.-L., Zhang, Y.-Z., et al. 2023b, *RAA*, 23, 085012
- Zhang, X., & Duan, R. 2022, *Proc. SPIE*, 12190, 1219032

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

Source: ChinaXiv — Machine translation. Verify with original.