Research on ROACH2-GPU Cluster Correlator–Design and Implementation of the F-engine Module (Postprint)

Authors: Niu Chenhui, Wang Qunxiong, Zheng Xiaoping, Tian Haijun, Wu Fengquan, Li Jixia, Chen Xuelei, Gao Jie

Date: 2017-09-26T00:00:00+00:00

Abstract

Correlators play a crucial role in radio astronomy. Traditional correlators predominantly employed Field-Programmable Gate Array (FPGA) or Application Specific Integrated Circuit (ASIC) technologies, characterized by lengthy development cycles and limited scalability for expansion and upgrades. In recent years, many newly developed radio interferometer array correlators have adopted architectures based on general-purpose FPGAs and Graphics Processing Units (GPUs). To address the requirements of the Dark Energy Radio Exploration Experiment (Tianlai Project), a heterogeneous correlator based on Reconfigurable Open Architecture Computing Hardware (ROACH2) and GPUs was developed. This design separates functions such as data acquisition and Fast Fourier Transform (FFT) from complex multiply-accumulate operations, thereby fully exploiting the hardware resources of FPGAs and the computational speed of GPUs. The correlator offers excellent scalability, with computational load distributable across different nodes according to actual processing capabilities, providing exceptional flexibility. It has been successfully deployed in the Tianlai Project.

Full Text

Preamble

Vol. 14 No. 1 Jan. 2017
Astronomical Research & Technology
ISSN 1672-7673

Research on the ROACH2-GPU Cluster Correlator
–Design and Implementation of the F-engine Module

Authors:
Niu Chenhui[1], Wang Qunxiong[1], Wu Fengquan[3], Li Jixia[1], Chen Xuelei[1], Hao Jie[3], Zheng Xiaoping[2], Tian Haijun[1],[2]

Affiliations:
1. National Astronomical Observatories, Chinese Academy of Sciences, Beijing
2. Central China Normal University, Wuhan, Hubei
3. China Three Gorges University, Yichang, Hubei
4. Institute of Automation, Chinese Academy of Sciences, Beijing

---

Abstract

Correlators play a crucial role in radio astronomy. Traditional correlators typically employ Field-Programmable Gate Array (FPGA) or Application-Specific Integrated Circuit (ASIC) technologies, which suffer from long development cycles and limited extensibility. In recent years, many newly developed radio interferometer arrays have adopted correlators built on generic architectures combining FPGA and Graphics Processing Unit (GPU). To meet the requirements of the radio detection experiment for dark energy (the Tianlai project), we have developed a heterogeneous correlator based on Reconfigurable Open Architecture Computing Hardware (ROACH2) and GPU. This design separates data acquisition and Fast Fourier Transform (FFT) functionality from complex multiply-accumulate operations, fully leveraging FPGA hardware resources and GPU computational speed. The correlator offers excellent scalability, with computational load distributable across different nodes according to actual processing capabilities, providing great flexibility. It has already been deployed in the Tianlai project.

Keywords: correlator; Fast Fourier Transform; ROACH2; Tianlai project

---

Main Text

Introduction

In radio astronomy, correlators are primarily used to perform correlation operations on signals from different antennas to obtain visibility data, which contains radio interferometric information that can be used to reconstruct sky maps [1]. Correlation is a fundamental signal processing operation that describes the degree of correlation between two signals at different times. The processing workflow can be divided into two categories: one approach computes the correlation function first and then performs a complex Fourier transform to the frequency domain; the other performs the opposite operation—first transforming signals to the frequency domain and then conducting complex multiplication. Since frequency-domain multiplication is equivalent to time-domain convolution, both approaches yield identical results. However, as the number of correlation paths N increases, the former requires N^2 steps while the latter only requires N steps [2], making the frequency-domain approach significantly more efficient.

With rapid advancements in radio technology, interferometer arrays are becoming increasingly large, with more antenna elements, imposing substantial demands on computational capability. For instance, the Tianlai project [3] currently has dozens of dual-polarization array elements in its first phase, with plans to expand to nearly a thousand elements in the second phase. Similarly, the international Square Kilometre Array (SKA) project will feature approximately 200 dish antennas in its first-phase mid-frequency array, with potential future expansion to several thousand elements. These developments demand ever-higher extensibility, portability, and computational performance from correlators.

Currently, most interferometer arrays employ ASIC and FPGA hardware for their correlators. Such systems are typically designed for specific arrays and purposes, making them difficult to port to other radio telescope arrays. Moreover, telescope upgrades often necessitate complete redevelopment, and modifying ASIC designs is particularly challenging. To address these requirements, we have developed a ROACH2-GPU architecture-based correlator that separates analog signal digitization via Analog-to-Digital Converters (ADC) and FFT operations from complex multiply-accumulate computations. The F-engine handles data acquisition and FFT, while the X-engine performs complex multiplication and accumulation of different frequency channels to obtain visibility data for final storage. This architecture offers good scalability—additional hardware can be incorporated as the antenna array expands—and flexibly distributes computational load according to processing capabilities.

The F-engine in this correlator is implemented using ROACH2, a platform developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) at UC Berkeley, which provides extensive development libraries. Numerous interferometer arrays internationally have adopted this platform for correlator development, including the Large-Aperture Experiment to Detect the Dark Ages (LEDA) and the Precision Array for Probing the Epoch of Reionization (PAPER) [4]. This platform has gained wide acceptance in the radio astronomy community due to its portability—models can be conveniently adapted to different interferometer arrays with varying sampling frequencies or FFT lengths.

Considering that the FPGAs used for signal acquisition already possess substantial computational resources, our design adopts a ROACH2-GPU heterogeneous approach to fully utilize these resources while addressing the intensive computational demands of large arrays. The FFT tasks are assigned to the front-end FPGA, while the computationally intensive but algorithmically simple multiply-accumulate operations are offloaded to high-performance GPUs, thereby enhancing overall correlator performance. This paper focuses on the F-engine module design and presents the overall architecture and test results.

1 F-engine Architecture Design

The F-engine is developed based on the ROACH2 platform, which provides a graphical visual programming design environment [5] with rich libraries including ADC interfaces, 10 GbE port mapping modules, and common FPGA algorithmic modules such as logic operations and FFT. The underlying library models are well-documented [6], facilitating the construction of radio astronomy backends like correlators and spectrometers. ROACH2 is an improved version of ROACH, offering enhanced functionality and hardware upgrades that increase versatility for radio astronomy applications.

In the correlator, the F-engine primarily digitizes analog signals through ADCs and performs FFT operations. It converts multi-channel signals into frequency channels and transmits the data to the X-engine. While the original PAPER correlator model [4] integrated all F-engine functionality within one or a few FPGAs, this approach, though functional, lacked flexibility. For better extensibility, we developed a more generic F-engine model that runs identical code on each ROACH2 board, with software-controlled parameters in the data packets to differentiate boards. This allows flexible concatenation of multiple F-engines to create correlators with arbitrary multiples of the base channel count.

The F-engine consists of four main modules: ADC input, polyphase filter bank (PFB), equalizer, and transpose, as illustrated in Figure 1. The computational load can be adjusted based on actual conditions by modifying parameters in the transpose and Ethernet modules to alter packet information, thereby adapting to backend X-engine processing nodes.

The ROACH2 platform maps FPGA registers to the Linux file system, enabling remote control via the Karoo Array Telescope Control Protocol (KATCP). This interface provides Python, C, and Ruby libraries for register read/write operations. We developed a control program for the F-engine that allows remote configuration and monitoring.

2 Implementation of ADC and FFT Functions in F-engine

The data acquisition function is implemented in the Input module, which primarily handles ADC sampling. The ADC module utilizes the CASPER library component ADC16x250_8, which supports 8-channel acquisition at 250 MHz sampling rate with 8-bit resolution on ROACH2. The ADC16x250_8 module outputs 8-bit data streams that feed into subsequent processing stages.

To suppress side lobe effects after FFT, a polyphase filter bank module is inserted before the FFT operation. This applies weighting values to each signal sample—multiplying by a window function such as Hanning or Hamming. The PFB module is implemented using the CASPER DSP library.

The FFT function employs complex FFT operations, where each FFT module processes two input signals as real and imaginary components, generating two output channels. This approach saves computation time and FPGA resources.

The FFT module produces a spectrum of 2,048 frequency points. With 32 input channels, 16 complex FFT modules suffice to process the ADC outputs. Each module first outputs the spectrum of one input channel, followed by the second channel' s spectrum.

3 Equalizer Implementation and Coefficient Range Determination

In typical receivers, frequency response varies across channels, which can affect results within limited dynamic ranges. To address this, an equalizer module applies a compensation process to FFT output data streams. Each frequency channel receives a 4-bit adjustment factor (EQ coefficient) to ensure uniform receiver response across all channels. The coefficient values are critical for reducing in-band ripple.

The coefficients are determined by inputting Gaussian white noise and calculating the root-mean-square (RMS) value of channel autocorrelation after N integrations:

$$RMS = \sqrt{\frac{\sum_{i=0}^{N-1} \text{Autocorr}[i]}{N}}$$

where Autocorr represents the autocorrelation value after N integrations. The equalizer output for each channel is a signed complex number with 18-bit real and imaginary parts. Based on CASPER empirical values, optimal linearity is achieved when the RMS range is 2-3; linearity degrades when RMS falls below 1 or exceeds 7.

During commissioning, the antenna is pointed to cold sky, and receiver gain is adjusted to set the input signal amplitude. Each channel' s coefficient is tuned until the autocorrelation RMS falls within the 2-3 range.

4 Transpose Function Implementation

After FFT, each channel' s data stream contains only its own spectrum. Cross-multiplication between different channels requires data from the same frequency point, necessitating a transpose operation that reorders data from channel-frequency to frequency-channel arrangement. This is also known as a corner-turn operation.

The transpose module accomplishes this reordering. As shown in Figure 3, the data flow is rearranged from channel-major to frequency-major order. ROACH2 is equipped with a 10 GbE network card providing four 10 GbE ports. After transposition, data is distributed across these ports.

The transpose module' s internal structure consists of first-level sub-modules (transpose0 through transpose3), each handling 1/4 of the frequency band. With an observation bandwidth of 125 MHz and FFT length of 1,024, the frequency resolution is 0.122 MHz. Each 10 GbE port transmits 1/4 of the frequency points (256 channels). Each first-level module contains four second-level submodules that reorder data streams by controlling read/write address shifting in dual-port RAM.

Figure 6 illustrates the address timing: write addresses increment sequentially but skip cycles to selectively write required frequency bands, while read addresses interleave to output data from two channels simultaneously, ensuring proper frequency-channel alignment.

5 Network Transmission Implementation

In the system network, each processing unit has its own MAC address. The module assigns destination IP addresses to each packet based on information from the transpose module. The 1,024 frequency channels are divided into four segments, each sent to a different X-engine node via a 10 GbE switch. Each packet contains 32 frequency points, with each point including data from all 32 channels of a single ROACH2 board. After equalization, each channel is truncated to 8 bits.

To verify network transmission integrity, an 8-byte CRC checksum is appended to each packet. The complete packet format is shown in Table 1. The packet header includes a timestamp counter for identifying transmission times and an FID (Frequency ID) to distinguish different ROACH2 boards and calculate corresponding channel numbers during correlation.

Two frequency resolution modes are supported: 0.122 MHz and 0.244 MHz. A RawData mode is also implemented for ADC verification, bypassing FFT and correlation to directly stream digitized waveforms to high-performance computers for storage and analysis.

6 Performance Testing

ADC Sampling Accuracy Test: The system includes a RawData packet mode that streams ADC samples directly to storage without processing. Testing with a 20 MHz sine wave input yielded a fitted frequency of 20.0001 MHz, consistent with theoretical expectations.

Correlator Linearity Test: Using passive sky as a signal source, the receiver gain was varied to control input amplitude while measuring channel autocorrelation outputs. As shown in Figure 8(b), the system demonstrates good linearity across an input range of -12 dBm to 6 dBm.

Interferometric Phase Test: Two experiments validated phase measurement accuracy. First, using FPGA-based delay modules with simulated noise, the measured phase slope was 1.561×10 rad/Hz, matching the theoretical value of 1.571×10 rad/Hz within 0.6% error (Figure 9). Second, using real noise with a 7.5 m coaxial cable length difference to create delay, the results confirmed expected phase-frequency relationships.

Astronomical Source Observation: The correlator was used with three cylindrical antennas and receivers to observe the Sun and Cygnus A. Offline processing clearly showed interferometric amplitude variations and phase fringes caused by source motion. Comparative observations with another correlator developed by the Institute of Automation, Chinese Academy of Sciences, showed consistent results [7].

7 Conclusion

The F-engine model and its control program successfully implement data acquisition, FFT, and network distribution functions. The correlator architecture offers easy scalability—additional hardware can be integrated when expanding the interferometer array without requiring new model development. The Raw_Data mode enables direct ADC data storage for verification.

As computational precision requirements increase, the F-engine' s workload may grow. However, with maturing GPU programming, future designs could offload FFT operations from FPGA to GPU backends, using ROACH2 primarily for data acquisition and network distribution. The ROACH2-based correlator' s extensibility and ease of development are gaining acceptance among radio astronomers, injecting new vitality into the field.

References

[1] Thompson A R, Moran J M, Swenson Jr G W. Interferometry and synthesis in radio astronomy. New York: John Wiley & Sons, 2001.

[2] Harris C. A parallel model for the heterogeneous computation of radio astronomy signal correlation. Australia: The University of Western Australia, 2010.

[3] Chen Xuelei. Radio detection of dark energy—Tianlai project. Scientia Sinica Physica, Mechanica & Astronomica, 2015, 45(12): 1358-1366.

[4] Bradley R F, et al. The precision array for probing the epoch of reionization. Astronomical Journal, 2010, 139(4): 1468-1480.

[5] Pei Xin, Li Jian, Chen Maozheng, et al. A design of a digital correlator for microwave holography based on the ROACH platform. Astronomical Research & Technology, 2015, 12(2): 219-227.

[6] Xu Yonghua, Li Jiyun, Zhang Yingqian, et al. A study of a pulsar observation system using the coherent de-dispersion technique. Astronomical Research & Technology, 2015, 12(2): 480-486.

[7] Wang Qunxiong, Niu Chenhui, Tian Haijun, et al. A research on the ROACH2-GPU-cluster-based correlator—the design and implementation of an X-engine model. Astronomical Research & Technology, 2017, 14(1): 219-227.

Author Information:

Niu Chenhui, Wang Qunxiong, Wu Fengquan, Li Jixia, Zheng Xiaoping, Tian Haijun, Chen Xuelei, Hao Jie

Abstract:

Correlator plays an important role in radio interferometry astronomy. Traditional correlator is achieved with FPGA or ASIC. It calls for long development cycle and its flexibility is unsatisfactory. At present, more and more radio astronomy arrays adopt generic architecture correlator with FPGA system and GPU. In order to meet the demands of the 'Tianlai' project, we developed a new correlator based on ROACH2 and GPU cluster. We divided the workflow of the correlator into two parts: we let the high performance GPU undertake the CMAC (conjugate multiply and accumulate) and only left the ADC and FFT for the F-engine. This correlator has high flexibility and it is easy to control. It has been applied to the Tianlai project.

Keywords: Correlator; F-engine; ROACH2; Tianlai project