中国科学院自动化研究所
INSTITUTE OF AUTOMATION
CHINESE ACADEMY OF SCIENCES

Digital Backend Solution for the Upgrade of Tianlai Pathfinder

21 cm Cosmology Workshop 2023

Liangtian Zhao, liangtian.zhao@ia.ac.cn

National ASIC Design Engineering Center
Institute of Automation, Chinese Academy of Sciences

2023.7.19

Outline

- **Digital Backend Solution for the Tianlai Pathfinder**
  - Overview of Tianlai Pathfinder
  - Upgrade plan for digital beamforming system
- **Digital backend devices of CASIA**
  - FPGA platforms(SNAP2、ChipK7、ChipRF)
  - Yellow blocks we have done
  - Module level products
  - Data processing servers
- **Applications based on CASIA devices**
  - Mingantu Meter-Decameter Radioheliograph Correlation System
  - 21 CentiMetre Array (21CMA) Digital Beamforming System
  - Qitai Radio Telescope Digital Backend

Overview of Tianlai Pathfinder Project

Classification	Antenna	Processing system	Observation target
Main array(A)	3 cylindrical antennas (40 meters in length, 15 meters in width in the north-south direction) with 96 feed sources, dual-polarization.	Beamforming System Correlator	FRB、Pulsar、 The 21cm spectral line of neutral hydrogen
Subarray(B、C)	Cylindrical antenna, with 32 feed sources	Beamforming System	FRB、Pulsar
Dish antenna Array	16 antennas with a diameter of 6 meters, Dual-polarization	Correlator Beamforming System	The 21cm spectral line of Classification、FRB、Pulsar

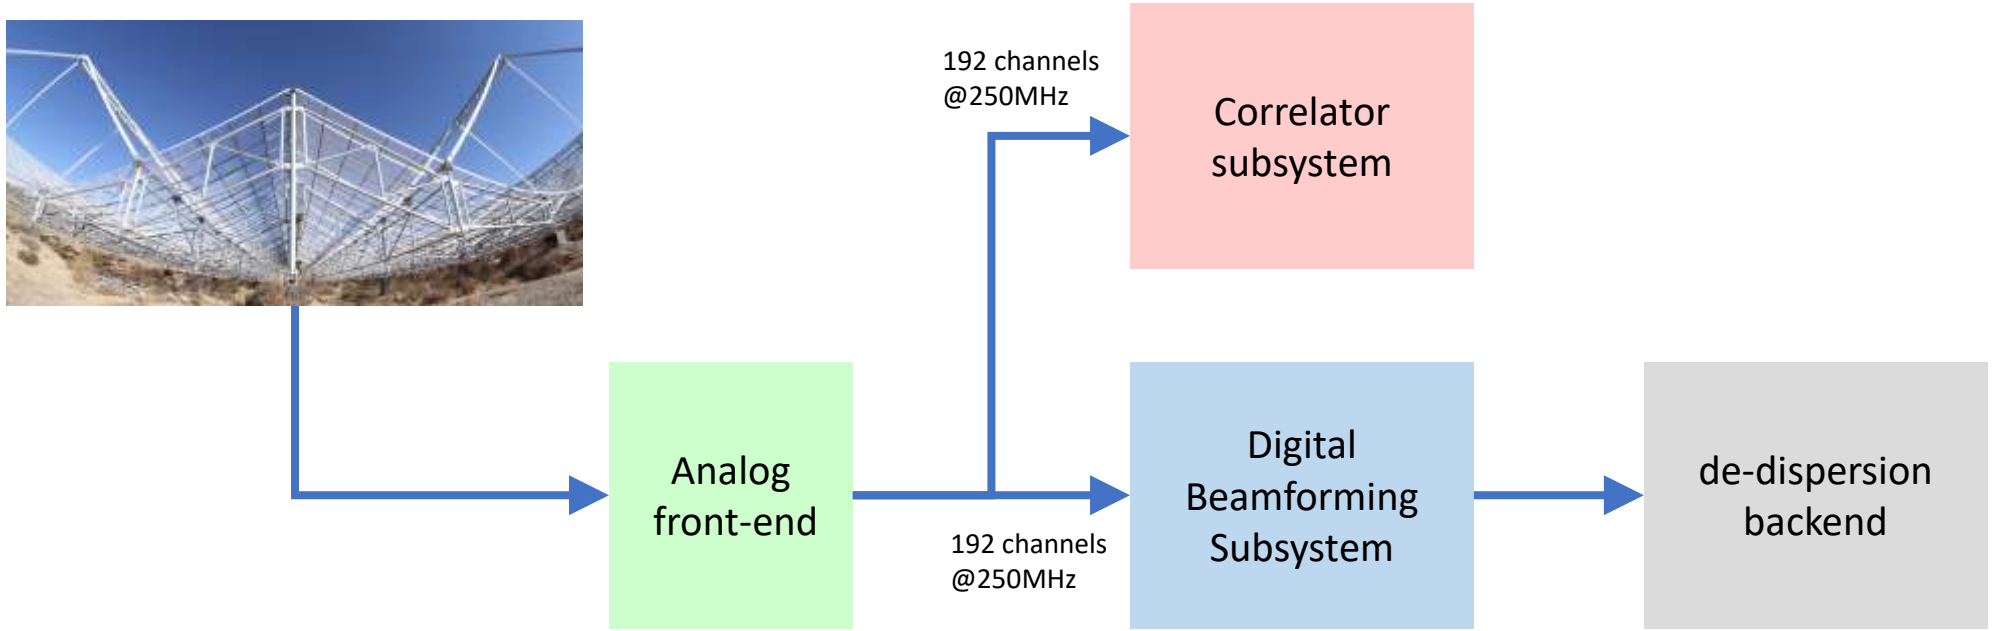

Tianlai Cylindrical Array Digital Beamforming System

System Architecture and Features

Architecture

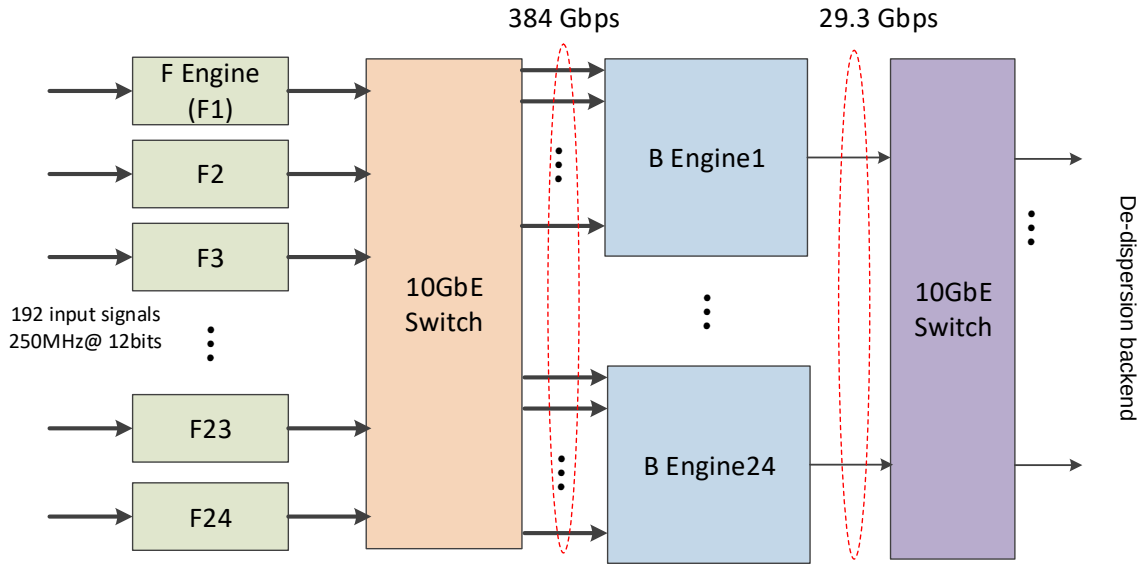

Features

Scientific goals	Features	Firmwares developed on toolflow(jasper)
Beamformer	IF inputs: **196** input channels IF Bandwidth: **DC~100MHz** Sample rate: **250Msps@12bits** FFT channels: 1024 Channel bandwidth: 122.1KHz Temporal resolution: **0.1ms** Quantification before beamforming: **8bit** Number of beams: **96**	Kintex7_channelization.slx

Beamformer Block Diagram

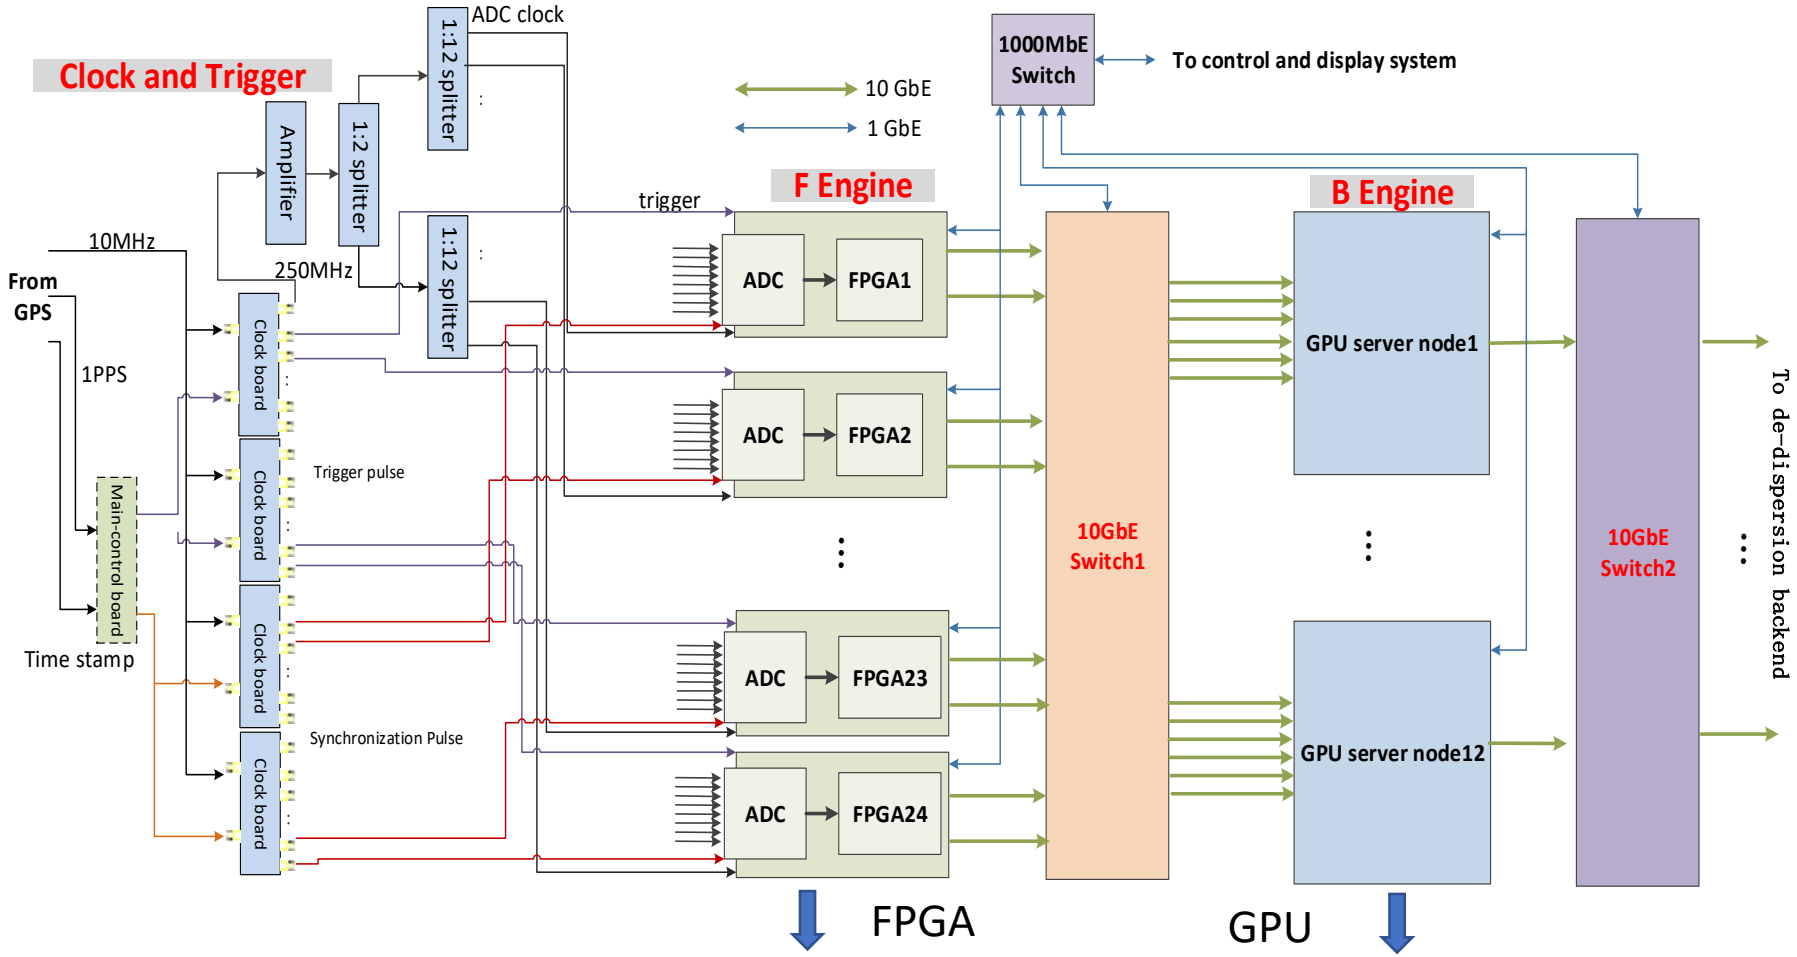

F-Engine

Signal Processing

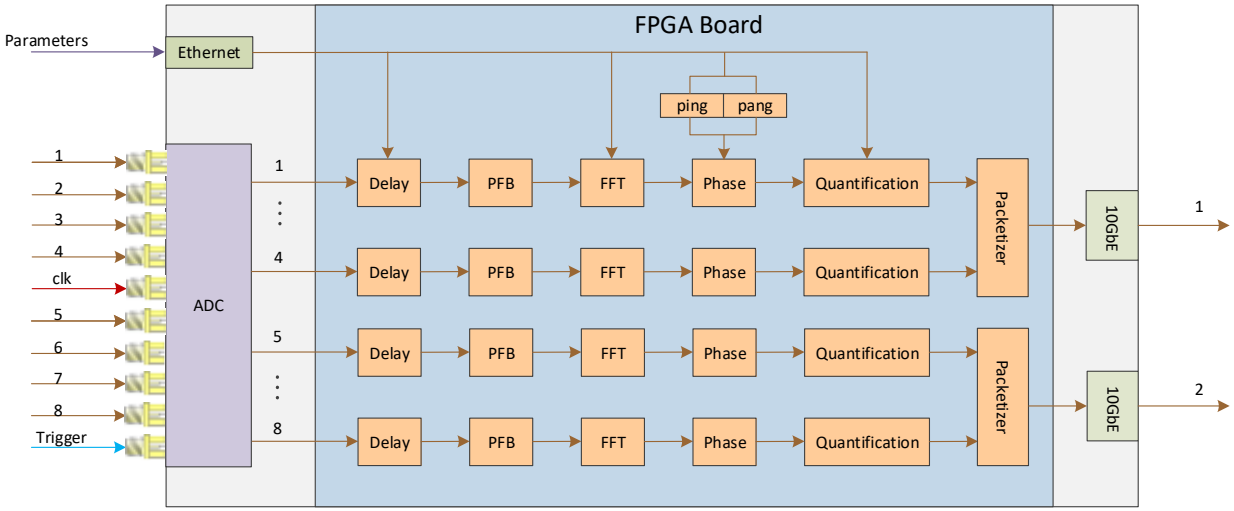

Hardware

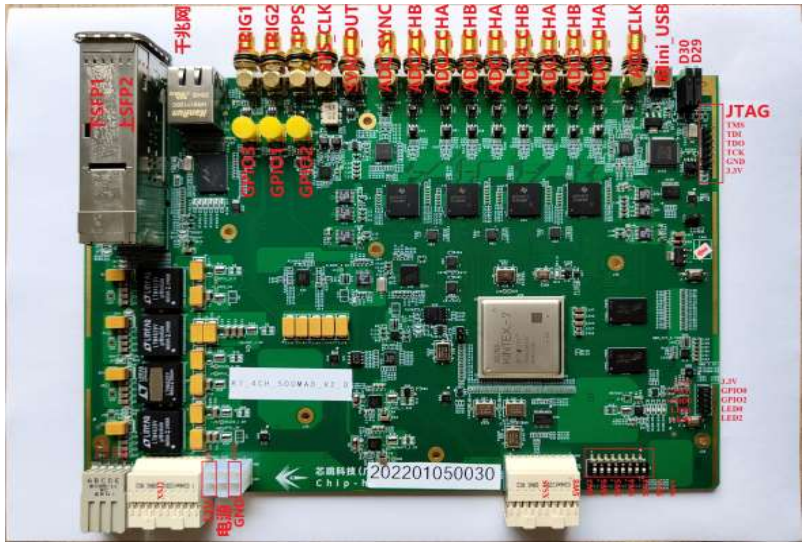

ChipK7, 8 inputs

10 x ChipK7, 80 inputs

B-Engine

- B-Engine consists of 24 CPU-GPU heterogeneous architecture servers.
- The CPU is responsible for data reception and result sending.
- The GPU is responsible for correlation, factoring and beamforming.

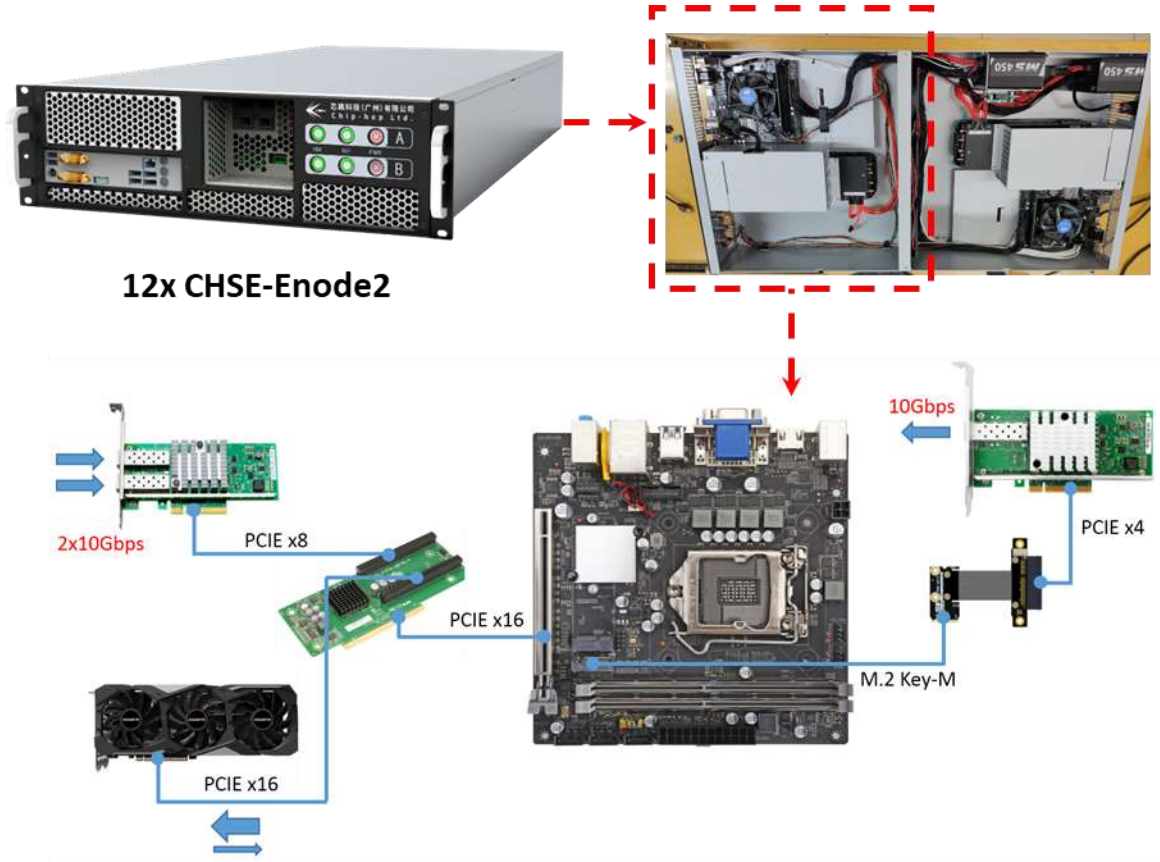

cost-effective server

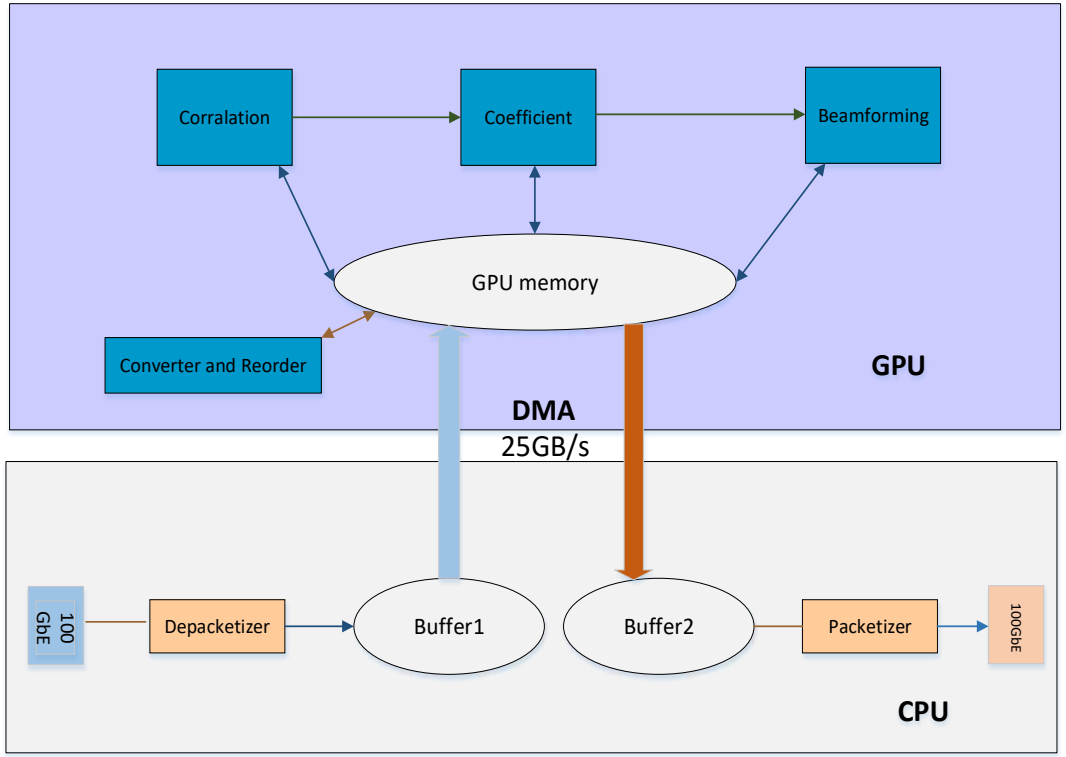

CPU-GPU design

Challenge: Synchronization

How to achieve synchronization between 192 input signals

- Synchronization of multiple ADCs

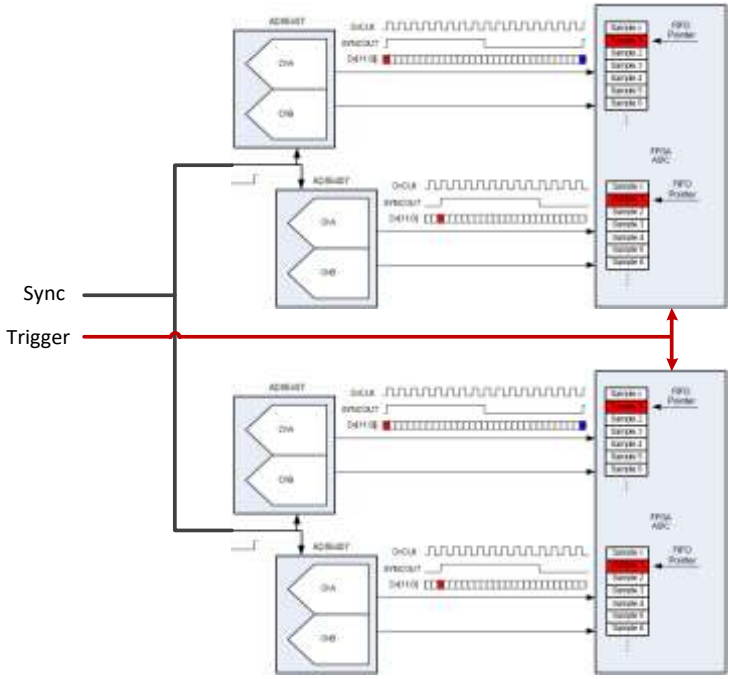

- Synchronization of multiple Boards

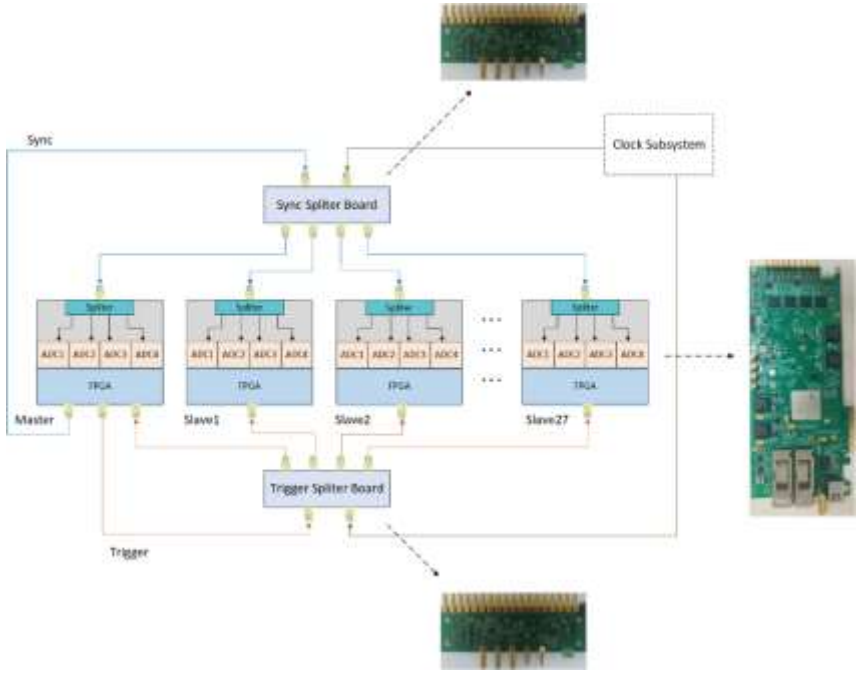

① Selecting ADC chip with synchronization function(TI ads5407)
② Equal length design of hardware circuit, especially for data and clock
③ Timing adjustment for data and clock using IDELAY in FPGA
④ Synchronized configuration, reset, trigger and data upload between boards by using multifunction pulses

Automatic synchronization after each power-on, synchronization accuracy <50ps

Challenge: Efficient data exchange

- **persistent high bandwidth** data exchange between F Engine and B Engine
  - MAC address binding
  - Regularized routing rules reduce network congestion

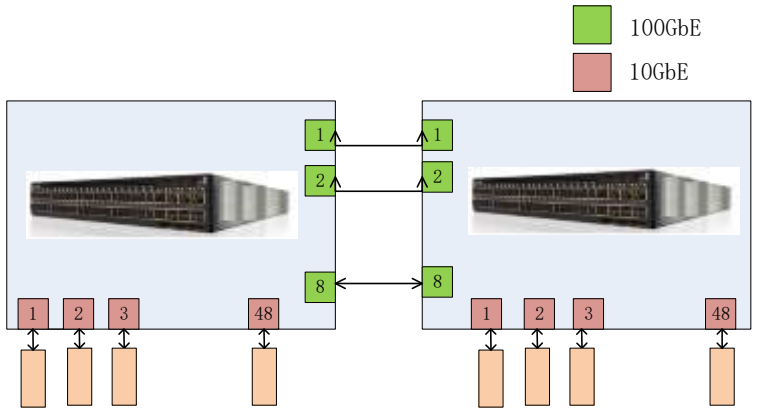

scalability

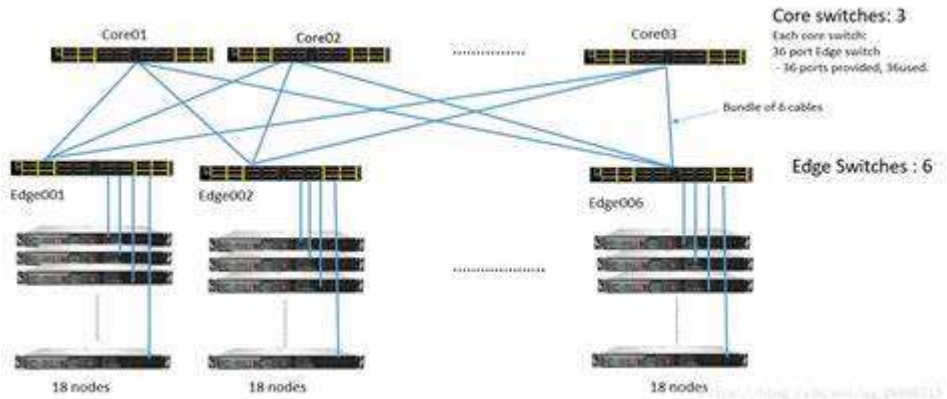

	Number of 10 Gbe ports	Bandwidth of each port
F Engine	48	8 Gbps
B Engine	48	8 Gbps

A fully interconnected network composed of multiple switches may have a decrease in effective bandwidth and have stability issues!

Challenge: Packet loss

Challenge： Under high-speed data bandwidth, data packets are easily lost, resulting in incomplete data, especially when the packet loss rate is too high, which affects the observation results.

- **Self developed high-speed network driver: User mode direct control reading, using polling instead of interrupt, achieving zero copy of data.**
- **Supports high and low speed bandwidth: 10GbE、40GbE、100GbE**

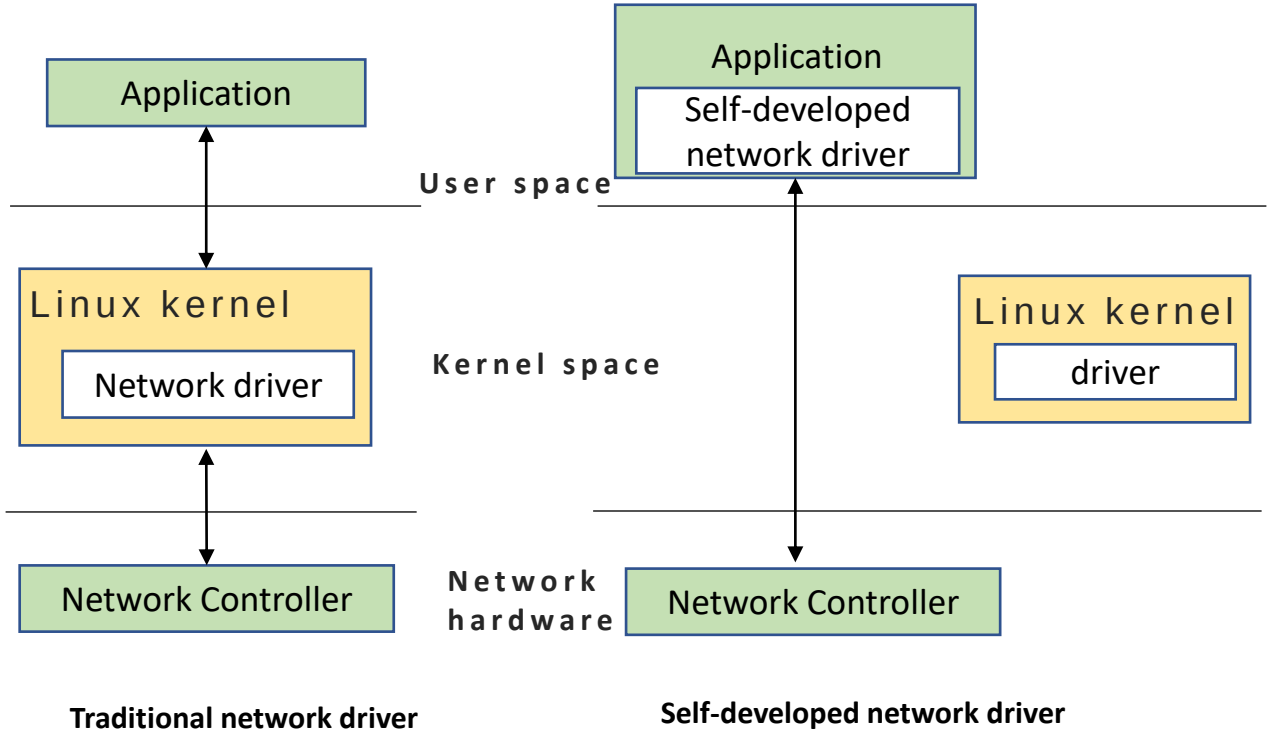

System Structure

@天籁
数字多波束合成子系统
V2.0

- Beam Form
- F Engine
- 资源监控
- 使用说明

运行状态

接收数据中

数据持续接收中

控制面板

连控制器　节点电源　连接状态　程序运行

电源管理

开启电源　关闭电源

程序控制

启动程序　停止程序

节点运行状态列表

IP	主机名	节点运行状态	备注
172.0.0.100	tianlai-172-0-0-100	接收数据中	
172.0.0.101	tianlai-172-0-0-101	接收数据中	
172.0.0.102	tianlai-172-0-0-102	接收数据中	
172.0.0.103	tianlai-172-0-0-103	接收数据中	
172.0.0.104	tianlai-172-0-0-104	接收数据中	
172.0.0.105	tianlai-172-0-0-105	接收数据中	
172.0.0.106	tianlai-172-0-0-106	接收数据中	
172.0.0.107	tianlai-172-0-0-107	接收数据中	
172.0.0.108	tianlai-172-0-0-108	接收数据中	
172.0.0.109	tianlai-172-0-0-109	接收数据中	
172.0.0.110	tianlai-172-0-0-110	接收数据中	
172.0.0.111	tianlai-172-0-0-111	接收数据中	
172.0.0.112	tianlai-172-0-0-112	接收数据中	
172.0.0.113	tianlai-172-0-0-113	接收数据中	
172.0.0.114	tianlai-172-0-0-114	接收数据中	
172.0.0.115	tianlai-172-0-0-115	接收数据中	

链路监控

未补偿的ADC | 补偿后的ADC | 未相位补偿的频谱 | 相位补偿后的频谱

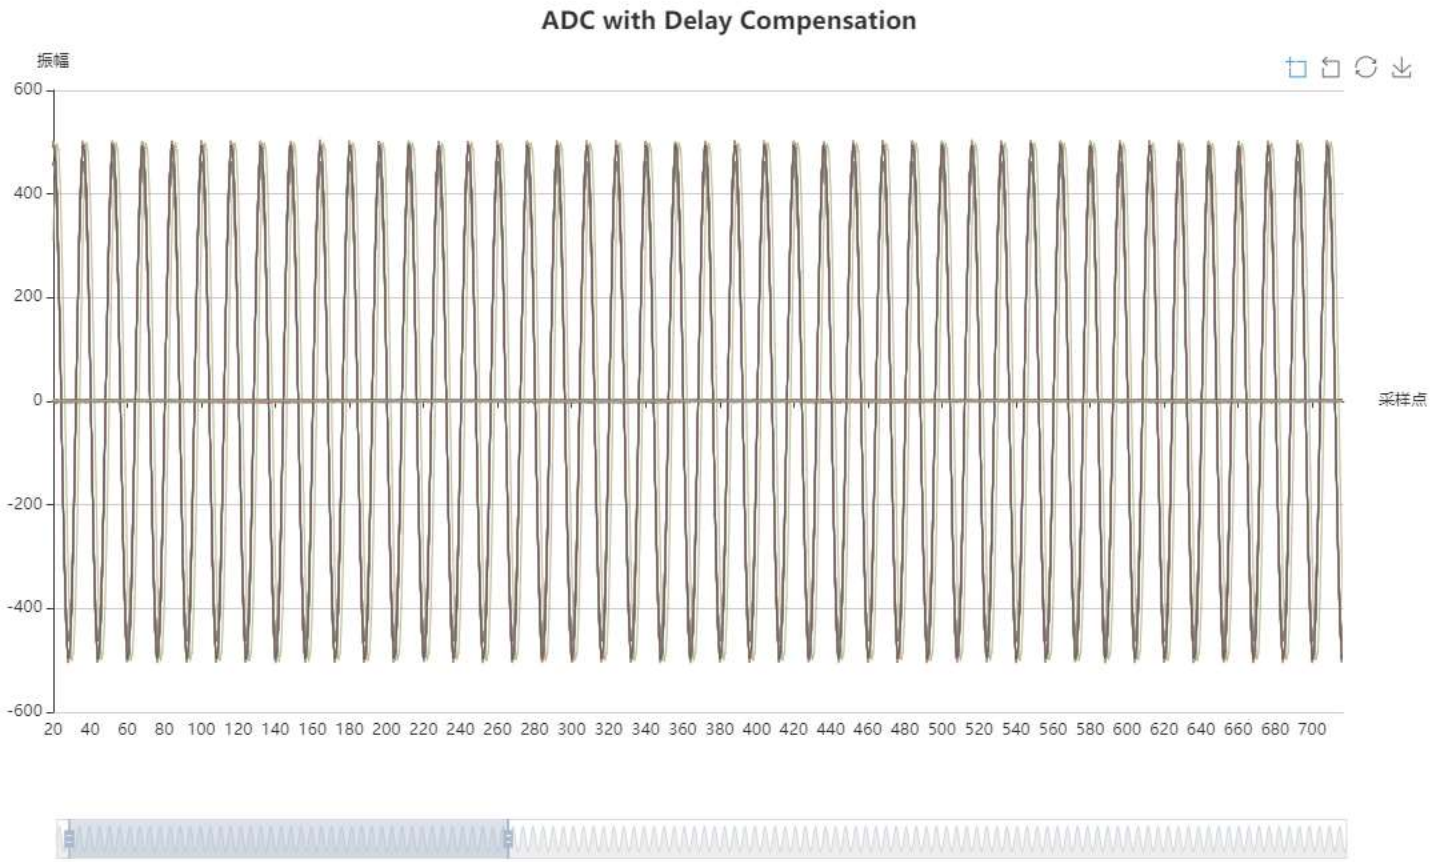

图形控制器

全选 反选

板卡 通道 + −

J&J

@天籁
数字多波束合成子系统
V2.0

- Beam Form
- F Engine
- 资源监控
- 使用说明

天籁数字多波束合成子系统

Last 6 hours | 10s

服务器资源总览表

IP	主机名	运行时间	内存	CPU核	5m负载	CPU使用率	内存使用率	分区使用率*	磁盘读取*	磁盘写入*	连接数	TCP_tw	下载带宽*	上传带宽*
localhost:9100	tianlai	5.85 week	125.56 GiB	128	0.63	0.08%	4.09%	9.62%	0 B/s	39.55 kB/s	123	12	2.07 Mb/s	2.34 Mb/s

各计算节点CPU占用率

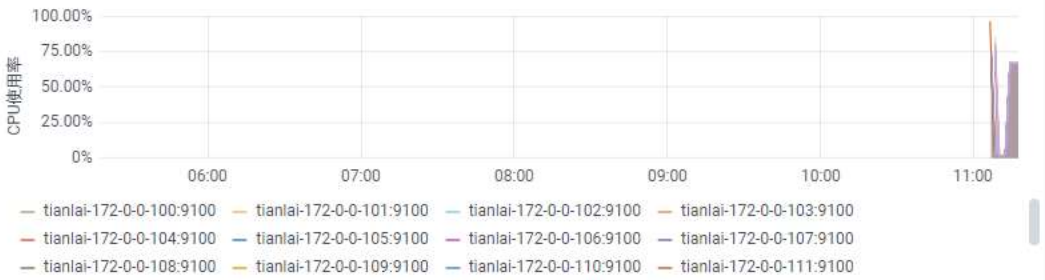

各计算节点内存使用率

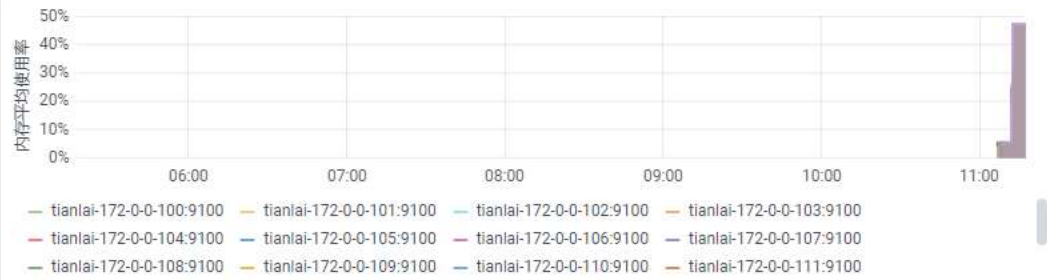

各GPU利用率

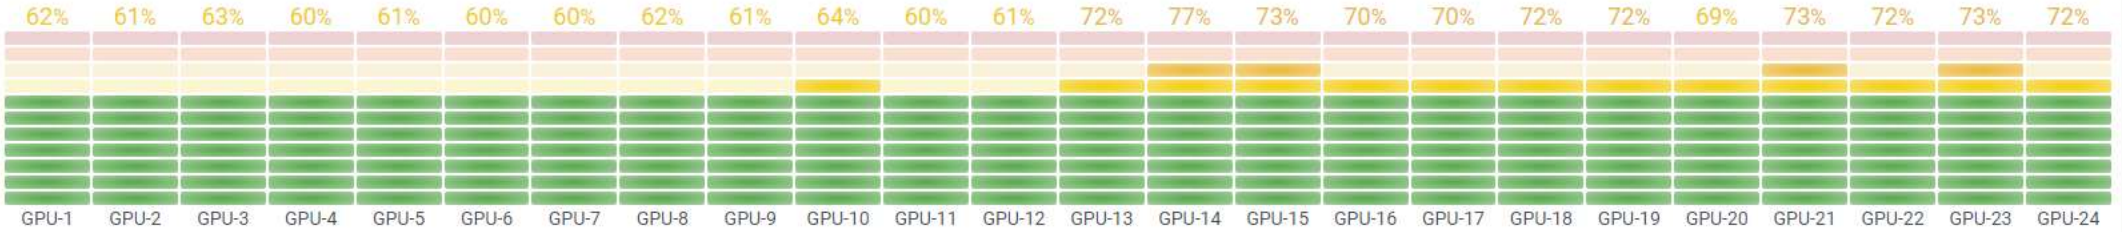

各GPU温度

各节点CPU核心温度

Upgrade Plan—Mainarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
96 feeds Single polarization	Input channels: 96 FPGA Boards: 12, 8 inputs each Output bandwidth: 192Gbps，8Gbps x 24 Port(10GbE)	Beams: 96 Computation: 4.66TFlops(Mul)+4.64Tflops(Add)=9.3Tflops GPU servers：RTX3060×12 Output bandwidth: 29.30Gsps	Input: 8Gbps, 24Port Output: 8Gbps, 24Port
		Beams: 192 Computation: 9.29TFlops(Mul)+9.26Tflops(Add)=18.55Tflops GPU servers：RTX3060×36；RTX4090×12 Output bandwidth: 58.59Gsps	Input: 8Gbps, 24Port Output: RTX 3060, 5.3Gbps, 36Port RTX 4090, 8Gbps, 24Port
		Beams: 288 Computation: 13.88TFlops(Mul)+13.92Tflops(Add)=27.8Tflops GPU servers：RTX3060×54；RTX4090×18 Output bandwidth: 87.9Gsps	Input: 8Gbps, 24Port Output: RTX 3060, 3.6Gbps, 54Port RTX 4090, 5.3Gbps, 36Port

Upgrade Plan—Mainarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
96 feeds Dual polarization	Input channels: 192 FPGA Boards: 24, 8 inputs each Output bandwidth: 384Gbps，8Gbps x 48 Port(10GbE)	Beams：96 per polarization，192 Computation: 9.31TFlops(Mul)+9.29Tflops(Add)=18.6Tflops GPU servers：RTX3060×24 Output bandwidth: 58.59	Input: 8Gbps, 48Port Output: 8Gbps, 48Port
		Beams: 192 per polarization，384 Computation: 18.58TFlops(Mul)+18.53Tflops(Add)=37.1Tflops GPU servers：RTX3060×72；RTX4090×24 Output bandwidth: 117.19Gsps	Input: 8Gbps, 48Port Output: RTX 3060, 5.3Gbps, 72Port RTX 4090, 8Gbps, 48Port
		Beams: 288 per polarization，576 Computation: 27.8TFlops(Mul)+27.77Tflops(Add)=55.61Tflops GPU servers：RTX3060×108；RTX4090×36 Output bandwidth: 175.78Gsps	Input: 8Gbps, 48Port Output: RTX 3060, 3.6Gbps, 108Port RTX 4090, 5.3Gbps, 72Port

Upgrade Plan—Mainarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
192 feeds Single polarization	Input channels: 192 FPGA Boards: 24, 8 inputs each Output bandwidth: 384Gbps，8Gbps x 48 Port(10GbE)	Beams: 96 Computation: 9.29TFlops(Mul)+9.28Tflops(Add)=18.56Tflops GPU servers：RTX3060×24 Output bandwidth: 29.3Gsps	Input: 8Gbps, 48Port Output: 8Gbps, 48Port
		Beams: 192 Computation: 18.53TFlops(Mul)+18.50Tflops(Add)=37.03Tflops GPU servers：RTX3060×72; RTX4090×24 Output bandwidth: 58.59Gsps	Input: 8Gbps, 48Port Output: RTX 3060, 5.3Gbps, 72Port RTX 4090, 8Gbps, 48Port
		Beams: 288 Computation: 27.77TFlops(Mul)+27.73Tflops(Add)=55.5Tflops GPU servers：RTX3060×108; RTX4090×36 Output bandwidth: 87.89Gsps	Input: 8Gbps, 48Port Output: RTX 3060, 3.6Gbps, 108Port RTX 4090, 5.3Gbps, 72Port

Upgrade Plan—Mainarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
192 feeds Dual polarization	Input channels: 384 FPGA Boards: 48, 8 inputs each Output bandwidth: 768Gbps，8Gbps x 96 Port(10GbE)	Beams：96 per polarization，192 Computation: 18.58TFlops(Mul)+18.55Tflops(Add)=37.13Tflops GPU servers：RTX3060×48 Output bandwidth: 58.59Gsps	Input: 8Gbps, 96Port Output: 8Gbps, 96Port
		Beams: 192 per polarization，384 Computation: 37.06TFlops(Mul)+37.0Tflops(Add)=74.06Tflops GPU servers：RTX3060×144; RTX4090×48 Output bandwidth: 117.19Gsps	Input: 8Gbps, 96Port Output: RTX 3060, 5.3Gbps, 144Port RTX 4090, 8Gbps, 96Port
		Beams: 288 per polarization，576 Computation: 55.54TFlops(Mul)+55.46Tflops(Add)=111.0Tflops GPU servers：RTX3060×216; RTX4090×72 Output bandwidth: 175.78Gsps	Input: 8Gbps, 48Port Output: RTX 3060, 3.6Gbps, 216Port RTX 4090, 5.3Gbps, 104Port

Upgrade Plan—Subarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
32 feeds Single polarization	Input channels: 32 FPGA Boards: 4, 8 inputs each Output bandwidth: 64Gbps， 8Gbps x 8 Port(10GbE)	Beams: 32 Computation: 0.53TFlops(Mul)+0.52Tflops(Add)=1.04Tflops GPU servers： RTX3060×4 Output bandwidth: 9.77Gsps	Input: 8Gbps, 8Port Output: 8Gbps, 8Port
		Beams: 64 Computation: 1.05TFlops(Mul)+1.02Tflops(Add)=2.07Tflops GPU servers： RTX3060×12; RTX4090×4 Output bandwidth: 19.53Gsps	Input: 8Gbps, 8Port Output: RTX 3060, 5.3Gbps, 12Port RTX 4090, 8Gbps, 8Port
		Beams: 96 Computation: 1.57TFlops(Mul)+1.53Tflops(Add)=3.1Tflops GPU servers： RTX3060×18; RTX4090×6 Output bandwidth: 29.3Gsps	Input: 8Gbps, 8Port Output: RTX 3060, 3.6Gbps, 18Port RTX 4090, 5.3Gbps, 12Port

Upgrade Plan—Subarray antennas

Sample rate: 250Msps@12bits; FFT channels: 1024; Channel bandwidth: 122.1KHz;
Quantification before beamforming : 8bit; Temporal resolution: 0.1ms

Upgrade Plan	F-Engine	B-engine	Switch F -> B
32 feeds Dual polarization	Input channels: 64 FPGA Boards: 8, 8 inputs each Output bandwidth: 128Gbps，8Gbps x 16 Port(10GbE)	Beams：32 per polarization，64 Computation: 1.06TFlops(Mul)+1.03Tflops(Add)=2.09Tflops GPU servers：RTX3060×8 Output bandwidth: 19.53Gsps	Input: 8Gbps, 16Port Output: 8Gbps, 16Port
		Beams: 64 per polarization，128 Computation: 2.1TFlops(Mul)+2.05Tflops(Add)=4.15Tflops GPU servers：RTX3060×24; RTX4090×8 Output bandwidth: 29.3Gsps	Input: 8Gbps, 16Port Output: RTX 3060, 5.3Gbps, 24Port RTX 4090, 8Gbps, 16Port
		Beams: 96 per polarization，192 Computation: 3.14TFlops(Mul)+3.06Tflops(Add)=6.2Tflops GPU servers：RTX3060×36; RTX4090×12 Output bandwidth: 58.59Gsps	Input: 8Gbps, 16Port Output: RTX 3060, 3.6Gbps, 36Port RTX 4090, 5.3Gbps, 24Port

Outline

- **Digital Backend Solution for the Tianlai Pathfinder**
  - Overview of Tianlai Pathfinder
  - Upgrade plan for digital beamforming system
- **Digital backend devices of CASIA**
  - FPGA platforms(SNAP2、ChipK7、ChipRF)
  - Yellow blocks we have done
  - Module level products
  - Data processing servers
- **Applications based on CASIA devices**
  - Mingantu Meter-Decameter Radioheliograph Correlation System
  - 21 CentiMetre Array (21CMA) Digital Beamforming System
  - Qitai Radio Telescope Digital Backend

SNAP2

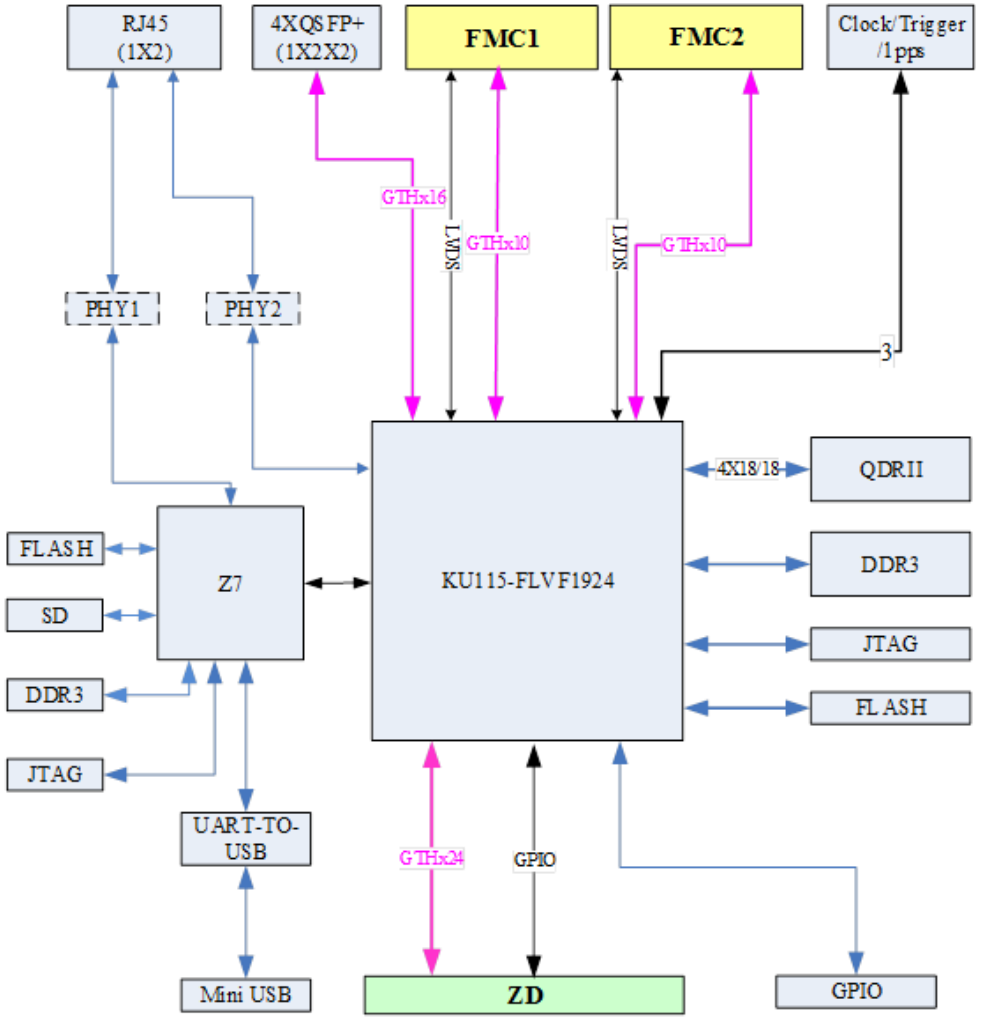

or

https://casper.berkeley.edu/wiki/SNAP2

http://www.chiphop.cn/ProductDetail/5143220.html

- **XCKU115**-FLVF1924 FPGA
- Zynq **XC7Z010**
- 4xQDRII
- 2xDDR3
- 4xQSFP+ connectors
- 2 Ethernet PHY interface with RJ-45 connectors
- 2xFMC HPC connectors
- 4xZD+ connectors
- 3xSMA: clock/1pps/trigger
- Micro secure digital (SD) connector
- 2xQuad spi flash
- Usb interface

reference
- Main Page
- Libraries
- Toolflow
- Hardware
- Software

documents
- Projects
- Tutorials
- Memos
- Papers
- Videos
- FAQ

wiki
- Recent changes
- Random page
- Help

search
Go Search

toolbox
- What links here

page discussion view source history

SNAP2

Background

SNAP 2 is a Kintex Ultrascale based platform, featuring a Xilinx XCKU115-FLVF1924 FPGA with 5520 DSP slices and 2160 36kb block

Design Features

- Ultrascale XCKU115-FLVF1924 FPGA.
- Zynq AP SoC XC7Z010
- 4 QDR memories
- 1 DDR3 component memory
- 4 Quad Small Form-factor Pluggable (QSFP) connectors, supporting 4x40GbE or 16x10GbE interfaces.
- 2 Ethernet PHYs supporting 1000BaseT
- USB interface
- ZD+ connector (this is not ZDOK!)
- 2 high pin count(HPC) FMC connectors

SNAP 2 is being designed by collaborators at the Institute of Automation, Chinese Academy of Sciences.

A more comprehensive overview of SNAP 2, including a block diagram can be found in File:SNAP2 Doc.pdf (direct link)

Update May 2016: The first prototypes of the SNAP2 board have been assembled! Pending testing by the Institute of Automation, one

Powerful Interface

FMC ADC Card(JESD204B, LVDS)

Make It Flexible !

ZD Expansion Board(10/40GbE, SRIO, PCIe x8)

FPGA Mezzanine Card(FMC)—ADC

- ADC1x3200-12
  - 1x TI ADC12D1600 **12bit** ADC
  - 2 inputs, **1.6 Gsps**
  - 1 input, **3.2 Gsps**

- ADC1x5000-10
  - 1x e2V EV10AQ190 **10bit** ADC
  - 4 inputs,1.25 Gsps
  - 2 inputs, 2.5 Gsps
  - 1 input, **5 Gsps**

- ADC8x250-8
  - 2x Hittite HMCAD1520 **8bit** ADC
  - 8 inputs**, 250 Msps**@8 bits
  - 4 inputs, 500 Msps @8 bits
  - 2 inputs, 1000 Msps @8 bits

- ADC1x6400-12
  - 1x TI ADC12DJ3200 **12bit** ADC
  - 2 inputs, 3.2 Gsps
  - 1 inputs, **6.4 Gsps**

ADC1x3200-12

ADC1x5000-10

ADC8x250-8

ADC1x6400-12

ChipK7

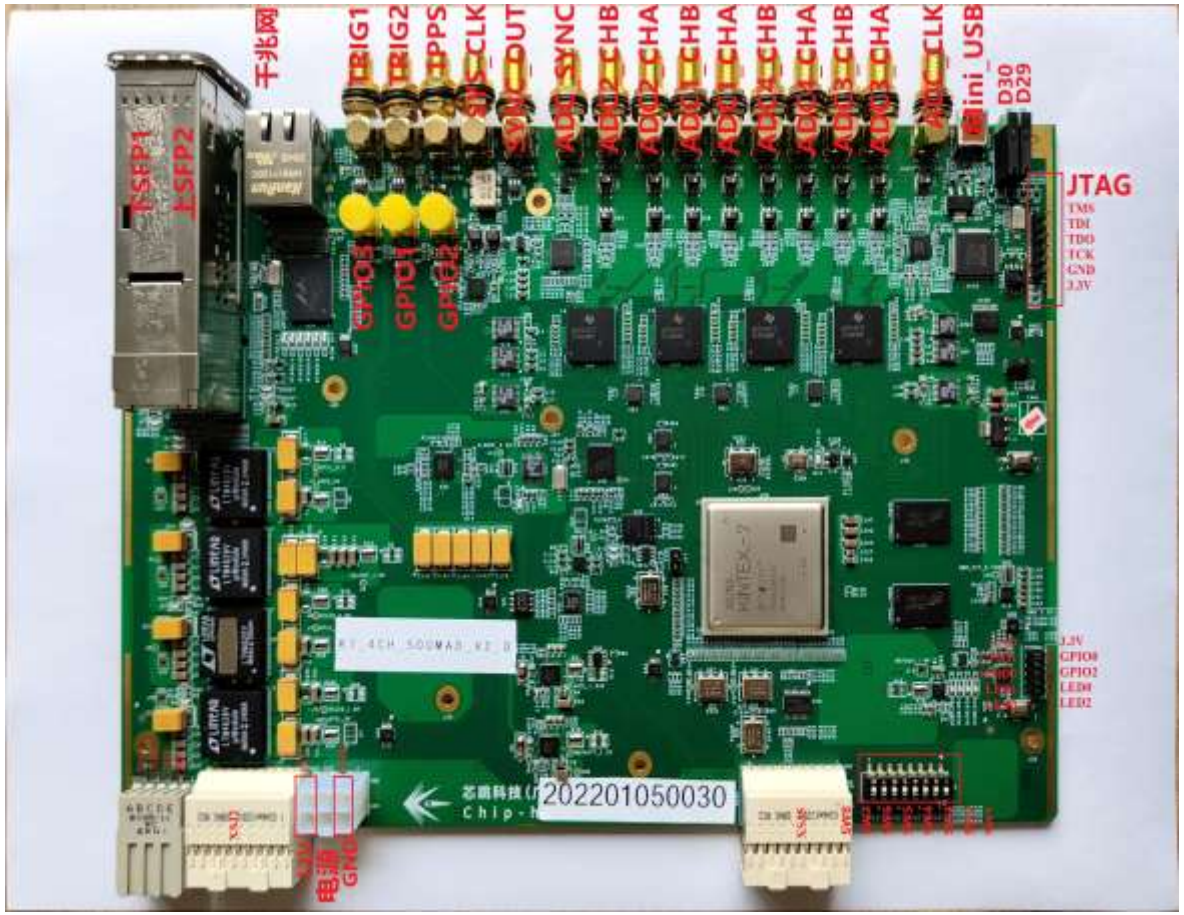

- KINTEX-7 XC7K325T FPGA
- x4 ADS5407 ADC chips (500Msps@12bit, 8 channel)
- x2 DDR3-2133×4Gb(32Meg×16bit×8Banks)
- x2 QSFP ports, supporting 8×10GbE interfaces
- x1 10/100/1000 Ethernet (RJ45) port
- x2 ZD+ ports: 16 x LVDS_GPIO/1PPS/SYNC/TRIG/ADC_CLK/SYS_CLK/Time Stamp
- x18 SMA: 8×ADC_CH/3×GPIO/2×SYNC/2×TRIG/1PPS/SYS_CLK/ADC_CLK
- x1 Mini USB interface for JTAG and serial port
- x4 LEDs
- x1 DIP switch

High performance and low-cost

ChipRF-UWL

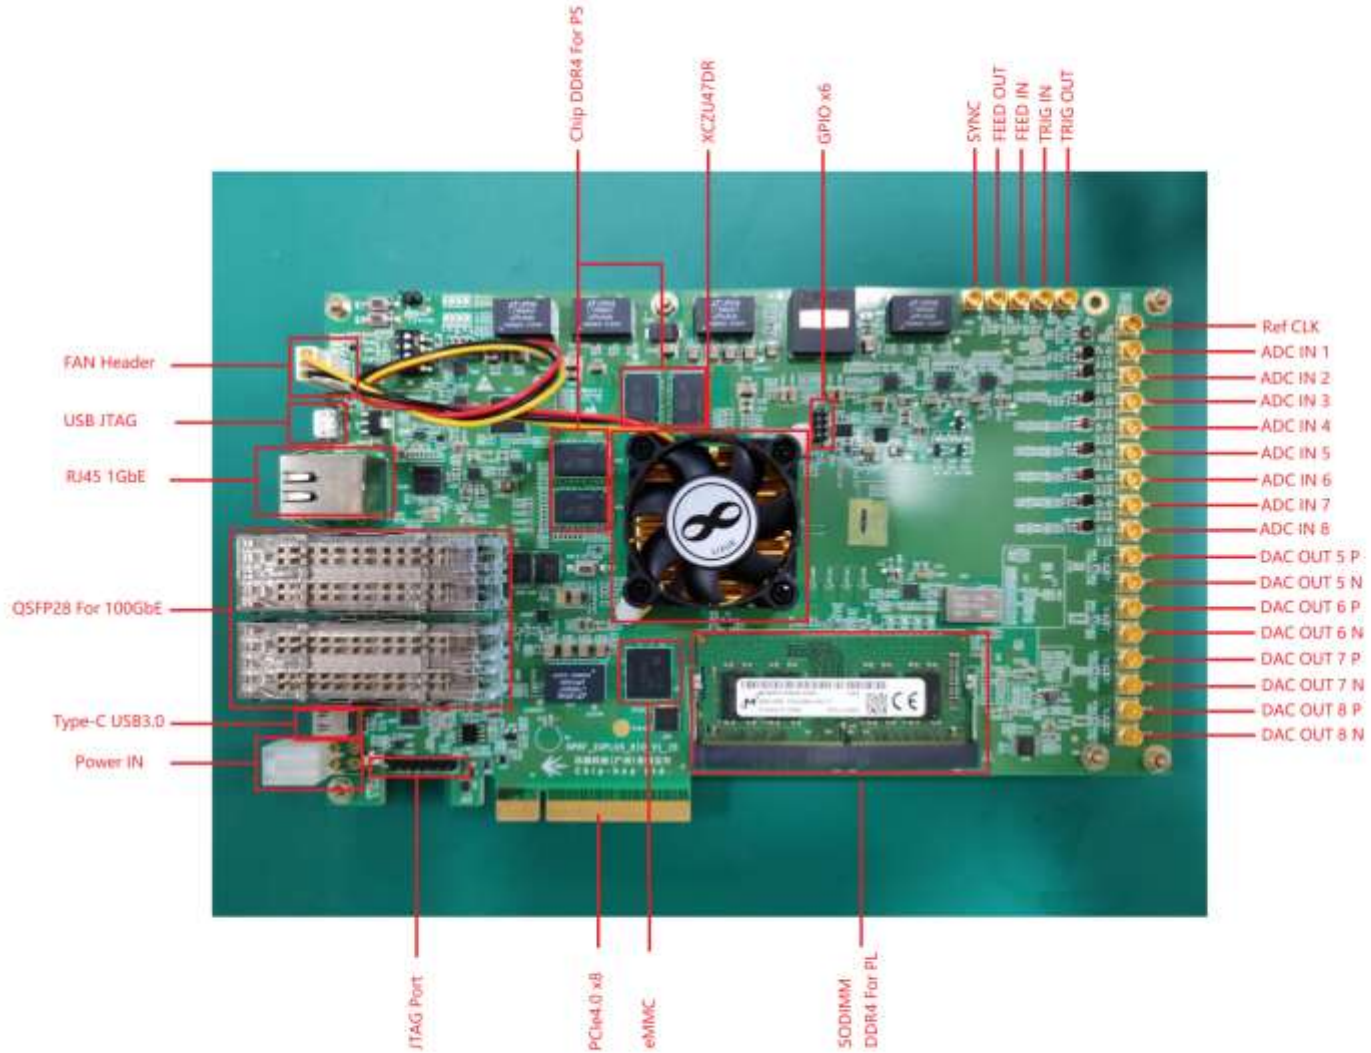

- ZYNQ UltraScale+ RFSoC XCZU47DR;
- x8 14bit ADC@5Gsps;
- x8 14bit DAC@9.85Gsps;
- x1 SODIMM for PL(8GB 2400Mb/s 64b);
- x4 DDR4 Chip for PS(4GB 2666Mb/s 64b);
- x2 QSFP28 ports, supporting 2×100GbE(CAUI-4/100GAUI-4) interfaces;
- x1 10/100/1000 Ethernet (RJ45) port
- x1 PCIe4.0x8 slot;
- x30 SMP: 8×ADC_CH, 16xDAC_DIFF, 1×REF_CLK, 1×SYNC, 2×TRIG_IN, 2×TRIG_OUT;
- x1 Mini USB interface for JTAG and serial port
- x4 LEDs, x6 GPIOs;
- X1 8bit eMMC(8~64GB);

RFSOC XCZU28DR/47DR; x8 14bit ADC@5Gsps

ChipRF-PAF

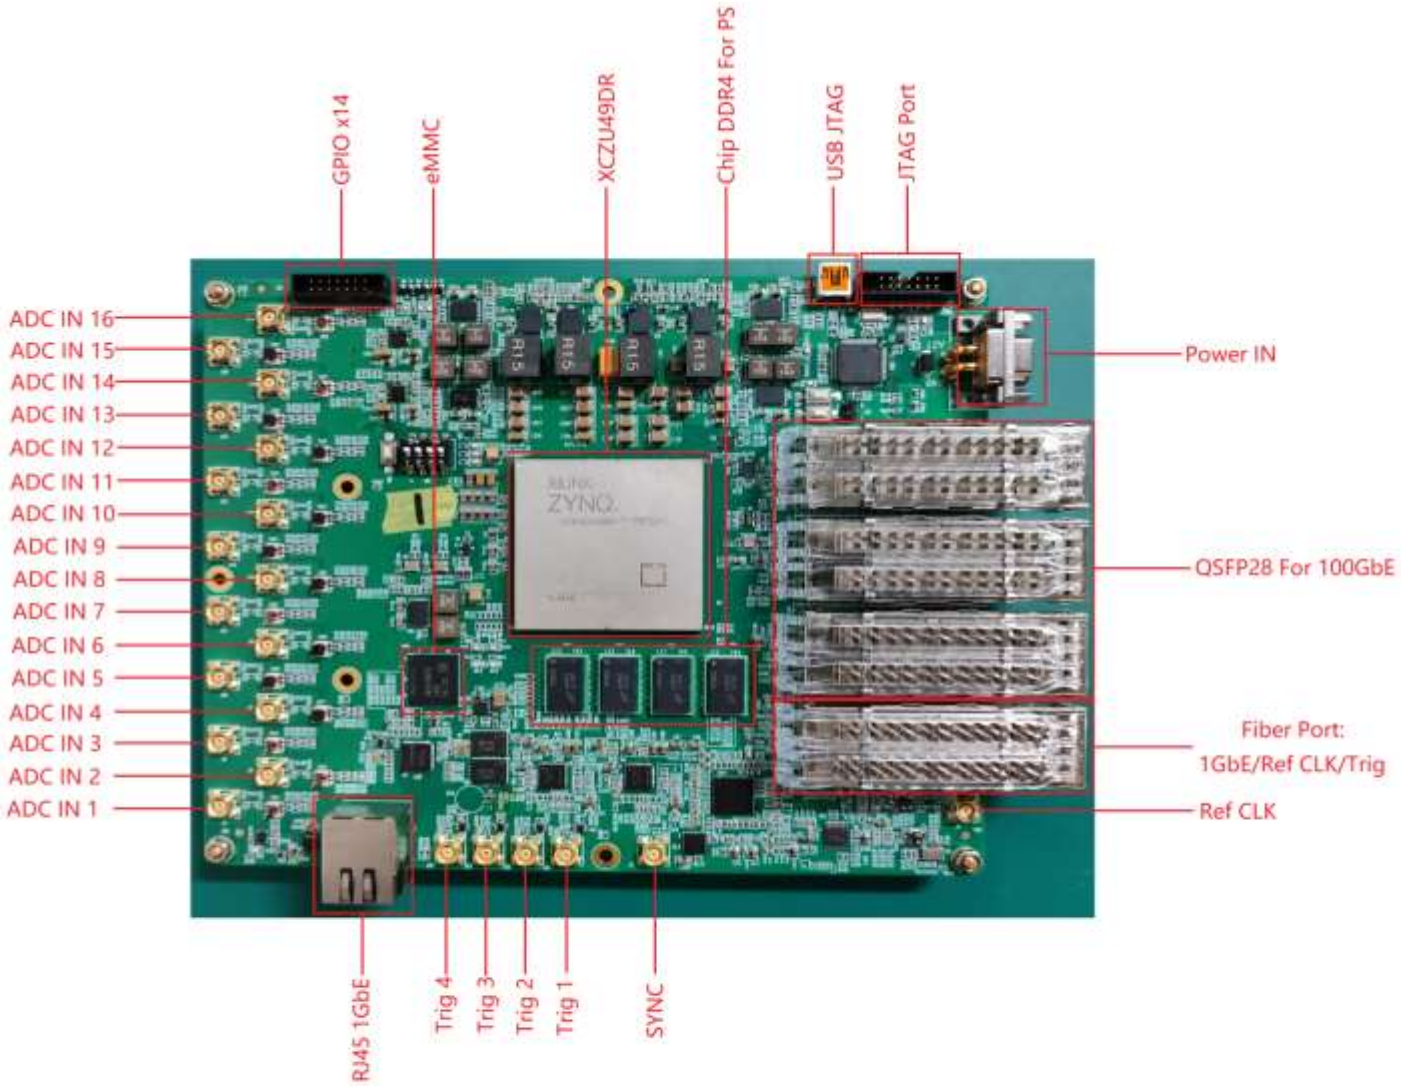

- ZYNQ UltraScale+ RFSoC XCZU49DR;
- x16 14bit ADC@2.5Gsps;
- x4 DDR4 Chip for PS(4GB 2666Mb/s 64b);
- x3 QSFP28 ports, supporting 3×100GbE(CAUI-4/100GAUI-4) interfaces;
- x1 10/100/1000 Ethernet (RJ45 or QSFP28) port
- x22 MMBX: 16×ADC_CH, 1×REF_CLK, 1×SYNC, 2×TRIG_IN, 2×TRIG_OUT;
- x1 Mini USB interface for JTAG and serial port
- x4 LEDs, x14 GPIOs;
- X1 8bit eMMC(8~64GB);

RFSOC XCZU49DR; x16 14bit ADC@2.5Gsps

Yellow Blocks We Have Done

- Three FPGA platforms have been ported to jasper toolflow

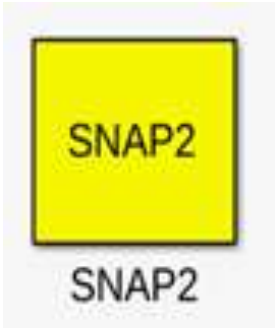

- Three kinds of ADC card have also been ported to jasper toolflow

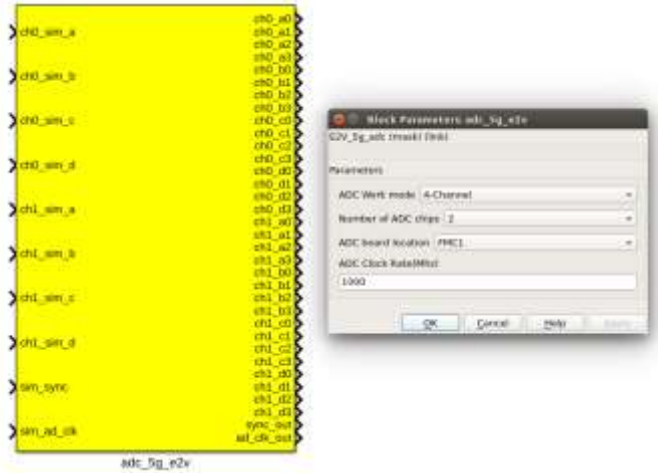

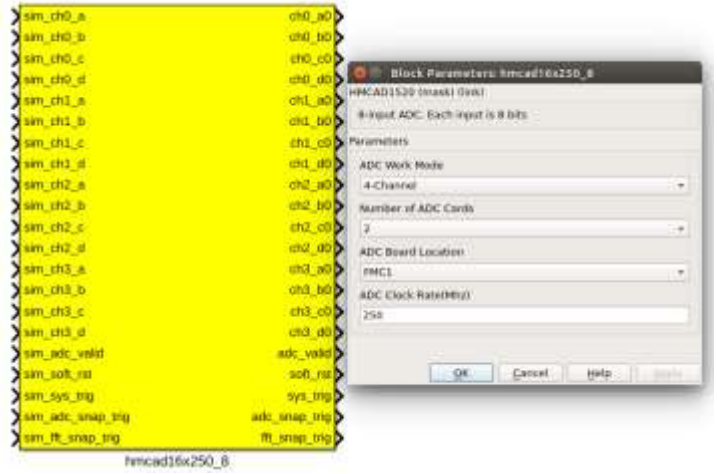

Links

https://github.com/shlean/mlib_devel or https://github.com/casper-astro/mlib_devel/tree/snap2

Module level products

Portable 8-channel Acquisition and Processing Module (CHRA-P500M)

Portable Ultra-wideband Acquisition and Processing Module (CHRA-PAD5G)

Low Latency SRIO Switch Module (CHRA-SR36)

Array Signal Synchronization Acquisition System (CHRA-DBF192)

Portable High-speed Data Recorder (CH-DR1000P)

Data processing server

High performance computing server
(CHSE-Slot8)

Computing accelerated expansion server
(CHSE-Slot16)

PCIe x16 cable

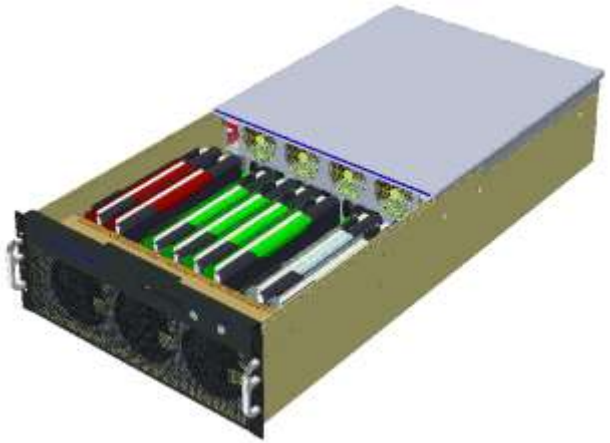

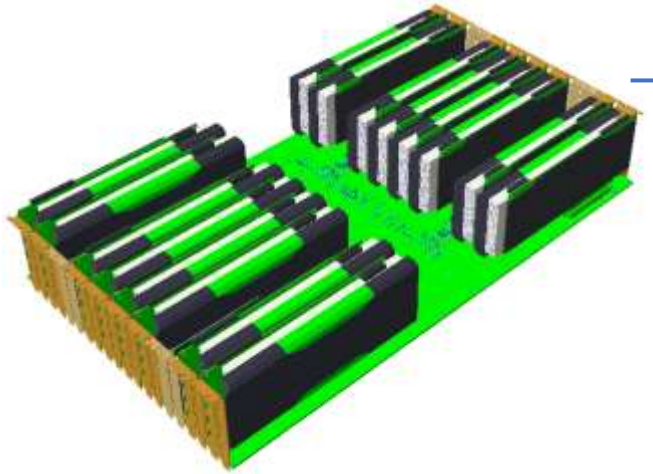

PCIe x16

- FPGA
- GPU
- NVMe SSD
- Smart NIC

Array signal acquisition and processing system

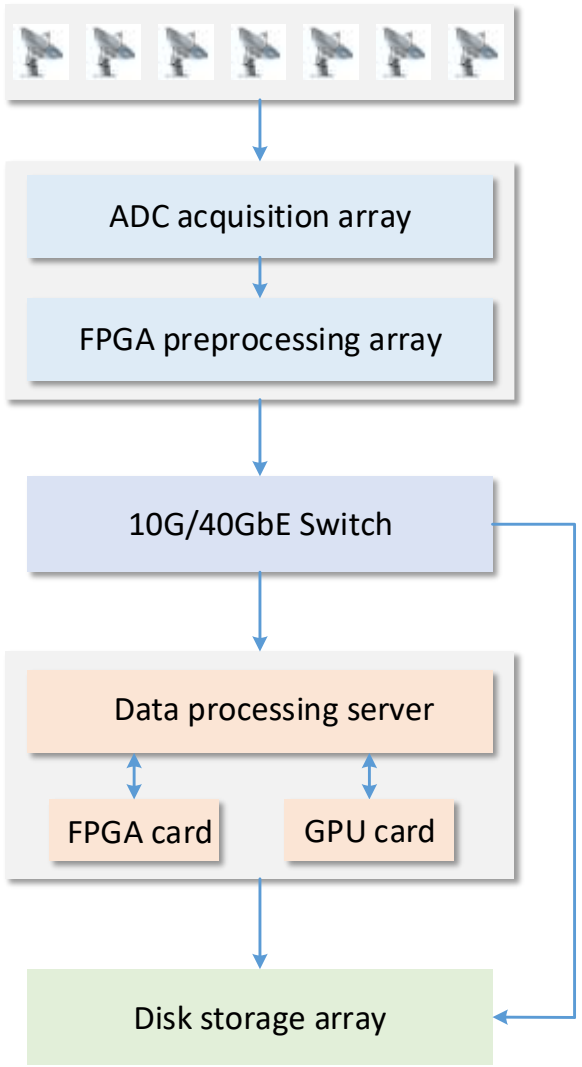

Display and Control

1GbE Switch

CHRA-DBF192

10GbE Switch

CHSE-Slot8

CHSE-Slot16

Disk Storage Array

Outline

- **Digital Backend Solution for the Tianlai Pathfinder**
  - Overview of Tianlai Pathfinder
  - Upgrade plan for digital beamforming system
- **Digital backend devices of CASIA**
  - FPGA platforms(SNAP2、ChipK7、ChipRF)
  - Yellow blocks we have done
  - Module level products
  - Data processing servers
- **Applications based on CASIA devices**
  - Mingantu Meter-Decameter Radioheliograph Correlation System
  - 21 CentiMetre Array (21CMA) Digital Beamforming System
  - Qitai Radio Telescope Digital Backend

Application:
MingantU Meter-Decameter Radioheliograph Correlation System

Mingantu, Inner Mongolia, China

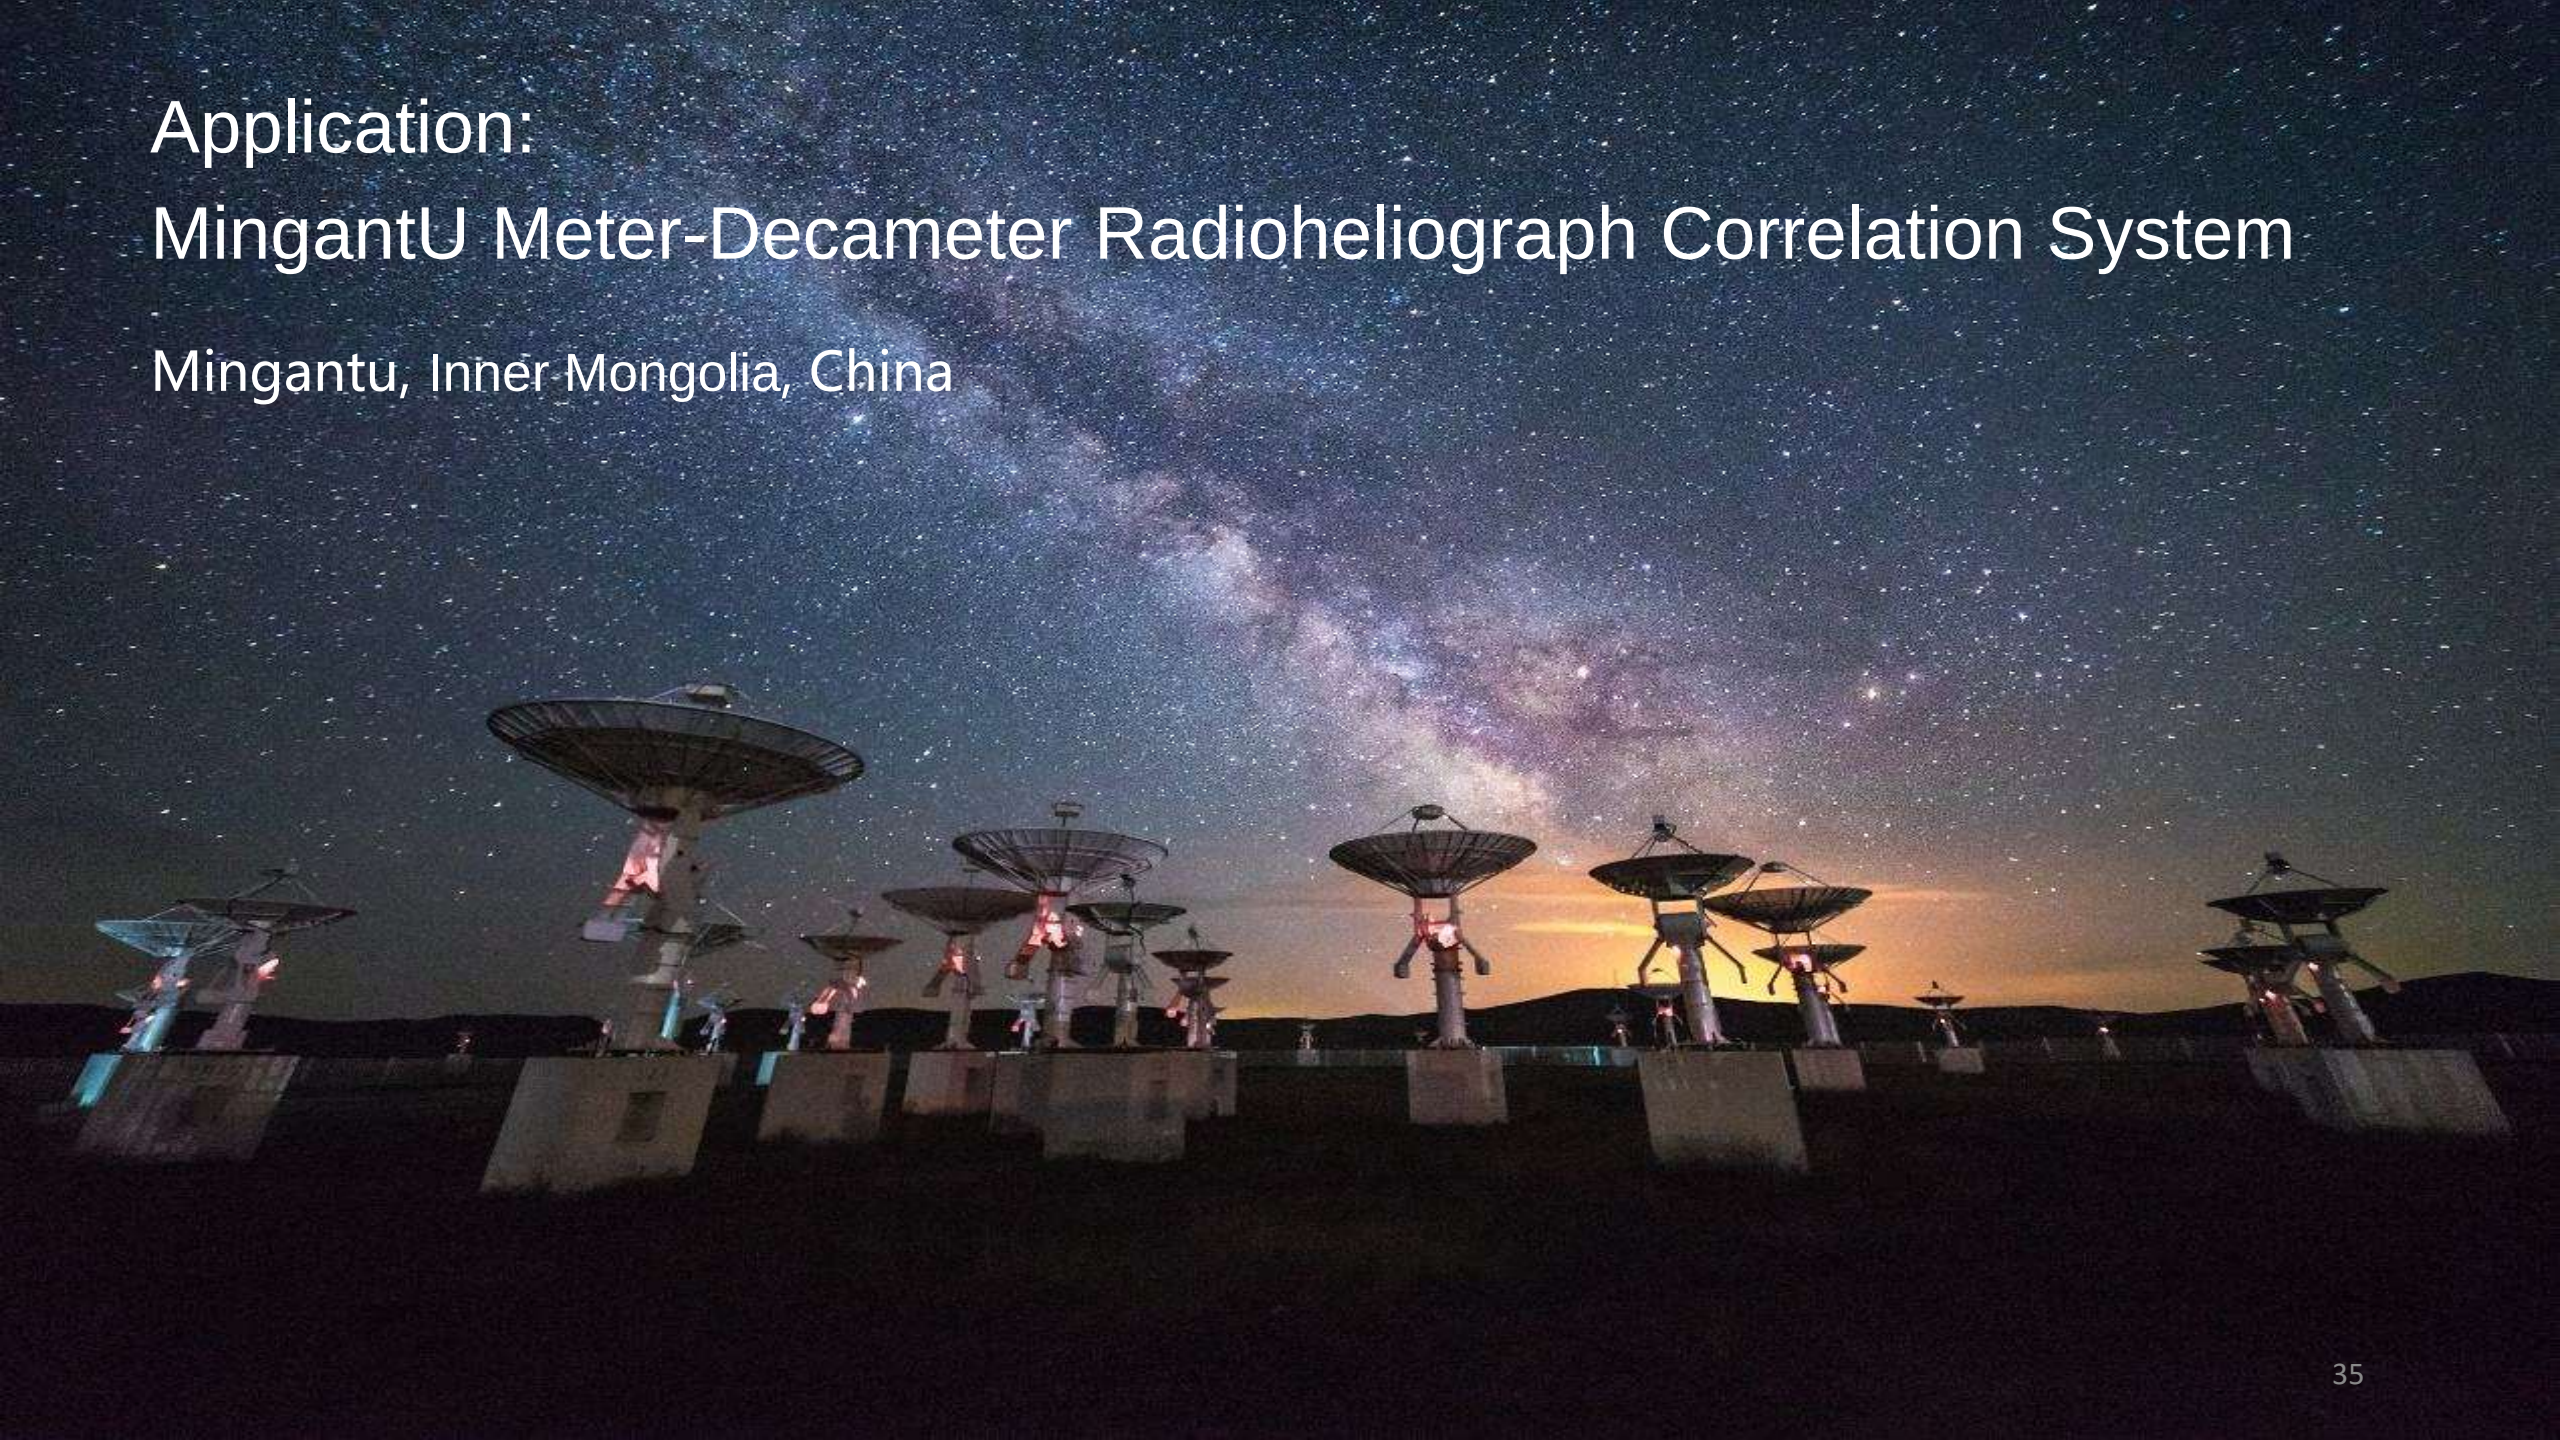

System Architecture and Features

Architecture

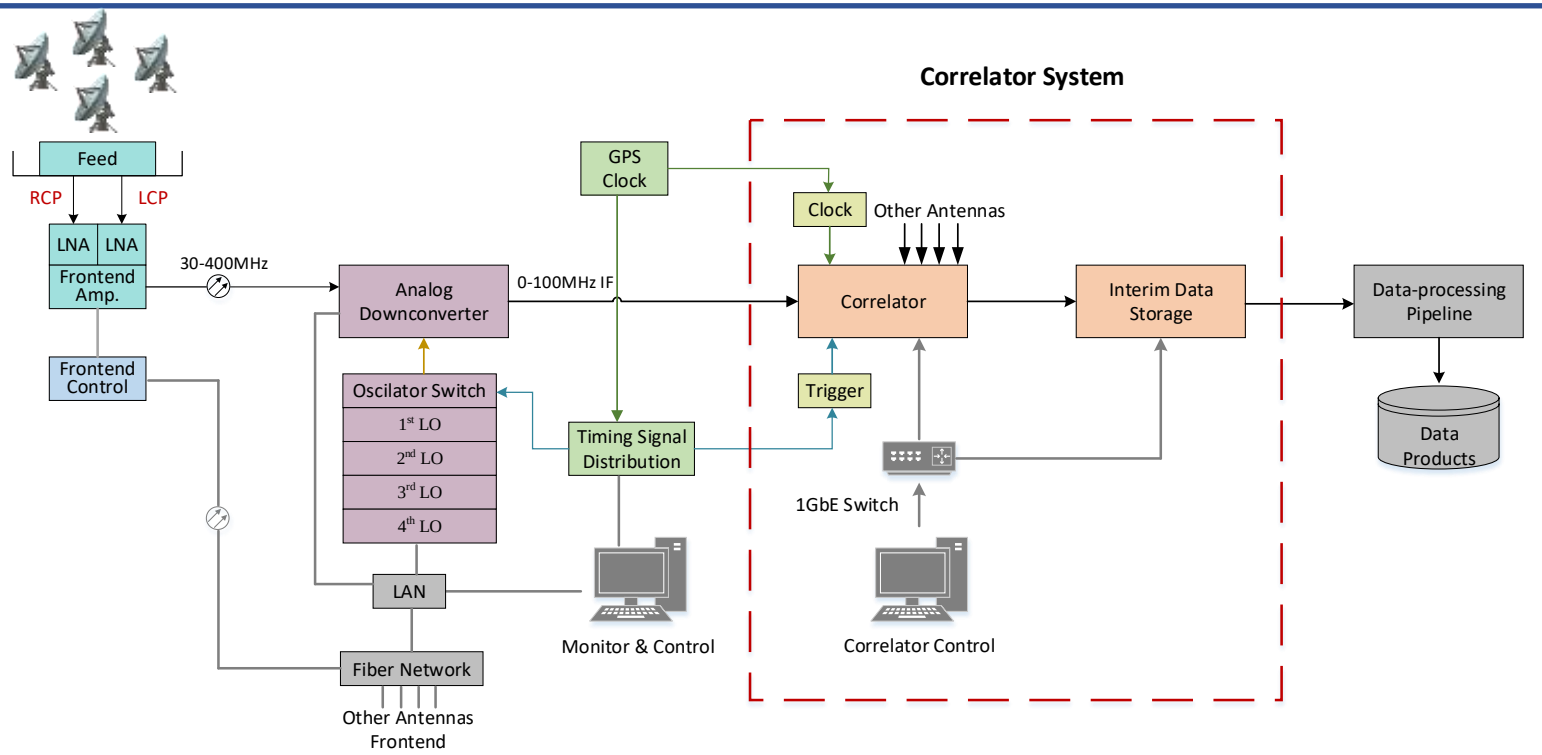

Features

Scientific goals	Features	Firmwares developed on toolflow(jasper)
Correlator	IF inputs: **101** antennas, Dual-polarization **202** input channels IF Bandwidth: **30~80MHz, 135MHz~235MHz** Sample rate: **250Msps@12bits**, bandpass sampling, synchronization FFT channels: **128** Channel bandwidth: 976.6KHz Temporal resolution: 25ms Quantification before correlation: 4bit	f_engine.slx x_engine.slx

Deployment

Thanks!

Email: liangtian.zhao@ia.ac.cn