

Design of VLBI2010 Software Correlator in IAA RAS

Igor Surkis, <u>Voytsekh Ken</u>, Alexey Melnikov, Vladimir Mishin, Dmitry Pavlov, Nadezhda Sokolova, Violet Shantyr, Vladimir Zimovsky

Institute of Applied Astronomy RAS, Russia

Introduction



XF correlator IAA RAS since 2009

- Input data stream of up to 1 Gb/s
- 6 stations/15 baseline
- 16 frequency channels
- Mark5B playback
- VSI-H format
- e-VLBI mode
- FPGA technology

UT 1-hour eVLBI sessions processing – daily EOP 24-hour sessions processing – weekly

DiFX since 2011

- Installed on a Sun Fire X4450
- Astrophysical experiments





Design goal is to provide the VLBI data processing for VLBI2010 small antenna network. The first two ones for the new national VLBI network are planned to be constructed by 2015. Software correlator is being developed in IAA RAS for nearrealtime processing.

Specifications:

- Input data stream of up to 16 Gb/s from each of up to 6 observatories:
 - > 2-bit sampling,
 - > 4 frequency bands:
 - 1 circlular polarization, 1024 MHz bandwidth, OR
 - 2 linear polarizations, 512 MHz bandwidth,
- Cross-spectra resolution of up to 4096 spectral channels
- Extracting of up to 16 phase calibration tones
- VDIF data format

Main design ideas



- The new correlator is an FX software one
- The basic principles are based on the correlator DiFX ideas
- The main distinctive feature is the use of Graphical Processing Units (GPUs) for the most computations, because GPU is equipped with hundreds of computing cores, and mathematical algorithms can be parallelized
- Hardware is based on the hybrid blade servers (CPU+GPU) in the high-performance computing cluster

Correlator equipment

High-perfomance computing cluster contains:

- GPU-based blade servers (approx. 20 units)
- Cache servers
- Cluster data storage
- Fiber optic based data transfer system
- Infiniband network (data rate of 56 Gb/s between any servers at a time)

Blade server structure



V200F compute module

- Up to 5 DP compute modules per chassis
- Over 14 TFLOPS in 5-node configuration (x86+GPU)
- 16 cores per compute module
- Two NVIDIA® Tesla™ M2090 GPUs
- 8 x 8 GB DDR3 RDIMMs (1066/1333/1600MHz)
- 2 cold swap disks (2.5", SAS/SATA 3Gb/s, up to 2TB)
- 2 external GbE ports and one optional FDR InfiniBand /40GbE VPI port
- Internal 100Mb/s management port (through the midplane)



The man in			
- (2010)			-
1.0. 100		<u> </u>	
12.382 (248 (248 (25	lit collect	R # 10	TARONE T



Correlator topology





SM – station module CM – correlation module

Station module





SSB – station signal block SM – station module



Decoding VDIF format ~ $16 \cdot 10^9$ operations = 16 GFLOPS Delay tracking ~ $2.1 \cdot 10^9$ operations = 2.1 GFLOPS SM contains 2 x Intel Xeon E5-2670 (8 cores @ 2.6 GHz) $2 \cdot 8$ (operations in tact) $\cdot 8$ (cores) $\cdot 2.6$ GHz = 332.8 GFLOPS

Buffering data, moving into/from RAM ~ 2x16 Gbps SM contains 8 x 8 GB DDR3 RDIMMs — more then 32 Gbps

Other operations — we have to make some experiments

So one station module provides all necessary operations for one station (we hope)

Correlation module



- Data transfer to DRAM GPU
- Transformation bit stream into float type
- Phase calibration signal extraction (up to 16 tones)
- Fringe rotation
- Auto- and cross-correlation computing with resolution of up to 4096 spectral channels



Data processing



- Fast pcal extraction method by S. Pogrebenko*
- Cross-correlation algorithm:



* - http://www.metsahovi.fi/~jwagner/pcal_extraction_practicals.pdf



 N=4096, 2 stations, 1 baseline, 1 frequency channel

 1 polarization, BW=1024 MHz
 T=480 (FFT)+24 (sum) = 504 GFLOPS

 2 polarizations, BW=512 MHz
 T=480 (FFT)+80 (sum) = 560 GFLOPS

 N=4096, 6 stations, 15 baselines, 4 frequency channels

 1 polarization, BW=1024 MHz each

 2 polarizations, BW=512 MHz each

 T=5760 (FFT)+672 (sum) = 6432 GFLOPS

 T=5760 (FFT)+2496 (sum) = 8256 GFLOPS

If we load the blade servers with efficiency of 25%, we need 1 correlation module for simplest correlator and 12 in case of six-station

IT'S A PRELIMINARY ESTIMATION!

Conclusion and plans



Present time – we know how to do it. We have the HPC cluster for expirements:

- 5 blade servers
- Data storage 24 TB
- Infiniband commutator

2013 yr – prototype programming, benchmarking and designing the one baseline one frequency channel correlator

2014-2015 yrs – 6-stations correlator design and start of processing





Thank you!