

Radioblocks GPU correlator

Aard Keimpema, Mark Kettenis





Image by Paul Boven (boven@jive.eu). Satellite image: Blue Marble Next Generation, courtesy of Nasa Visible Earth (visibleearth.nasa.gov).

Introduction

- Developed as part of the EU funded Radioblocks project
- Planned features
 - Reproduce the SFXC delay compensation
 - WOLA (Hann windowing)
 - Multiple phase centres
 - Coherent de-dispersion of transient science
- Timeline
 - Basic correlator: Finished in summer 2025
 - Advanced correlator: Deadline autumn 2026
- Goal of project is to asses feasibility of using GPUs for next generation EVN correlator

Grace Hopper



Grace CPU

CPU cores (Arm Neoverse V2)	72 Cores
Clock Frequency	3.1 GHz
LPDDR5X	Max. 480 GB
Memory bandwidth	Max. 512 GB/s
PCIe x16 (Gen 5) links	4x

Hopper H100 GPU

FP64	34 TFLOP
FP64 Tensor Core	67 TFLOP
FP32	67 TFLOP
TF32 Tensor Core	989 TFLOP
BFLOAT16 Tensor Core	1979 TFLOP
FP16 Tensor Core	1979 TFLOP
FP8 Tensor Core	3958 TFLOP
INT8 Tensor Core	3958 TFLOP
Power	450W - 1000W

Tensor cores have very high maximum performance


Tensor Core

$$\mathbf{D} = \begin{pmatrix} \begin{matrix} A_{0,0} & A_{0,1} & A_{0,\dots} & A_{0,15} \\ A_{1,0} & A_{1,1} & A_{1,\dots} & A_{1,15} \\ A_{\dots,0} & A_{\dots,1} & A_{\dots,\dots} & A_{\dots,15} \\ A_{15,0} & A_{15,1} & A_{15,\dots} & A_{15,15} \end{matrix} & \begin{matrix} B_{0,0} & B_{0,1} & B_{0,\dots} & B_{0,15} \\ B_{1,0} & B_{1,1} & B_{1,\dots} & B_{1,15} \\ B_{\dots,0} & B_{\dots,1} & B_{\dots,\dots} & B_{\dots,15} \\ B_{15,0} & B_{15,1} & B_{15,\dots} & B_{15,15} \end{matrix} \\ \text{FP16 or FP32} & \text{FP16} \end{pmatrix} + \begin{pmatrix} \begin{matrix} C_{0,0} & C_{0,1} & C_{0,\dots} & C_{0,15} \\ C_{1,0} & C_{1,1} & C_{1,\dots} & C_{1,15} \\ C_{\dots,0} & C_{\dots,1} & C_{\dots,\dots} & C_{\dots,15} \\ C_{15,0} & C_{15,1} & C_{15,\dots} & C_{15,15} \end{matrix} \\ \text{FP16 or FP32} \end{pmatrix}$$


- Accelerate matrix multiply-add operations
- Can handle mixed precision (accumulate in higher precision)
- Perfect match for (cross-)correlation
- Unfortunately, not delay compensation, coherent de-dispersion, FFT
- In VLBI correlation is ~30%-40% of computational load

VLBI RADIOBLOCKS

Prototype VLBI correlator



Bye Bye MPI; Welcome UCX




Abandoned MPI in favour of (lower-level) UCX

- Easier to deploy in heterogeneous environment
- Use wakeup feature to save energy (no polling)


Split correlator application into separate reader and correlator processes

- Created a “ucxwrappers” abstraction
- Can use RDMA directly into GPU memory
 - Even works for “localhost”; using ConnectX NIC as a DMA engine!
- Correlator process consumes almost no CPU cycles!
- Reader process doesn't need CUDA




Delay correction verification

Ef -O8





Phase VS Time

Ef -O8




Phase difference

Tensor Core correlator




• John Romein: <https://git.astron.nl/RD/tensor-core-correlator>

Tensor Core correlator




Significant speedup even in VLBI case



• John Romein: <https://git.astron.nl/RD/tensor-core-correlator>

Benchmark - Components



- Dummy correlation of RTX 3060
 - 32 Mhz bandwidth
 - 2 polarisations
 - 56s of “fake” data
- **Correlation is only 15% of total time**

Preliminary Benchmarks



EVN 16 stations, 512 MHz BW: 32.8 Gsamples/s, $O(2 \cdot 10^{-7})$ J/sample




Accelerator	Performance (Gsamples/s)	Efficiency (J/sample) *
RTX A4000 GPU	7.7	$1.8 \cdot 10^{-8}$
RTX 4000 Ada Generation GPU	9.3	$1.3 \cdot 10^{-8}$
A100 GPU	22.4	$1.1 \cdot 10^{-8}$
GH200 Grace Hopper Superchip	39.4	$2.5 \cdot 10^{-8}$ **

* Based on TDP/Maximum power

** Including CPU and memory

TODO: Coherent dedispersion


- Apply digital filter to remove dispersive delay from baseband signals
- Applied after delay correction but before cross-correlation
- Compute intensive
 - Very long FFTs, both forward and inverse
 - Overlap-add with half overlapping windows
- Reference implementation in SFXC



FRB 121102, DM = 561 pc.cm⁻³

TODO: Multiple phase centres

- Problem: Imaging the entire field of view yield very large data sets of mostly empty space
- Space is mostly empty, only few dozen sources in 10' field of view
- Solution: Correlate using high spectral / temporal resolution but only write small narrow field data sets for each source in the field.
- Probably can be integrated with Tensor Core Correlator (TCC)



Status

- RADIOBLOCKS GPU hardware at JIVE
 - AMD Epyc 9555P CPU: 64 cores
 - 256 GB Memory
 - 2 TB NVME storage
 - Four NVIDIA RTX PRO 4000 (Blackwell)
- Expected to be able to process 16+ stations at 4 Gbps in real-time
- Basic end-to-end correlation is working
- Coherent de-dispersion and Multiple phase centres under development
- Source code: <https://code.jive.nl/keimpema/sfxgpu>