Development of the Total Power GPU Spectrometer (TPGS) for the ALMA Wideband Sensitivity Upgrade

Jongsoo Kim
Korea Astronomy and Space Science Institute
On behalf of TPGS team

KASI: Jongsoo Kim (PI, PM),
Jongsuk Hong,
Won-Kee Park,
Heesun Yoon,
Yaocheng Chen

NAOJ: Takashi Nakamoto (PM),
Makoto Shizukami,
Isaac Goicovich,
Camilo Saldias

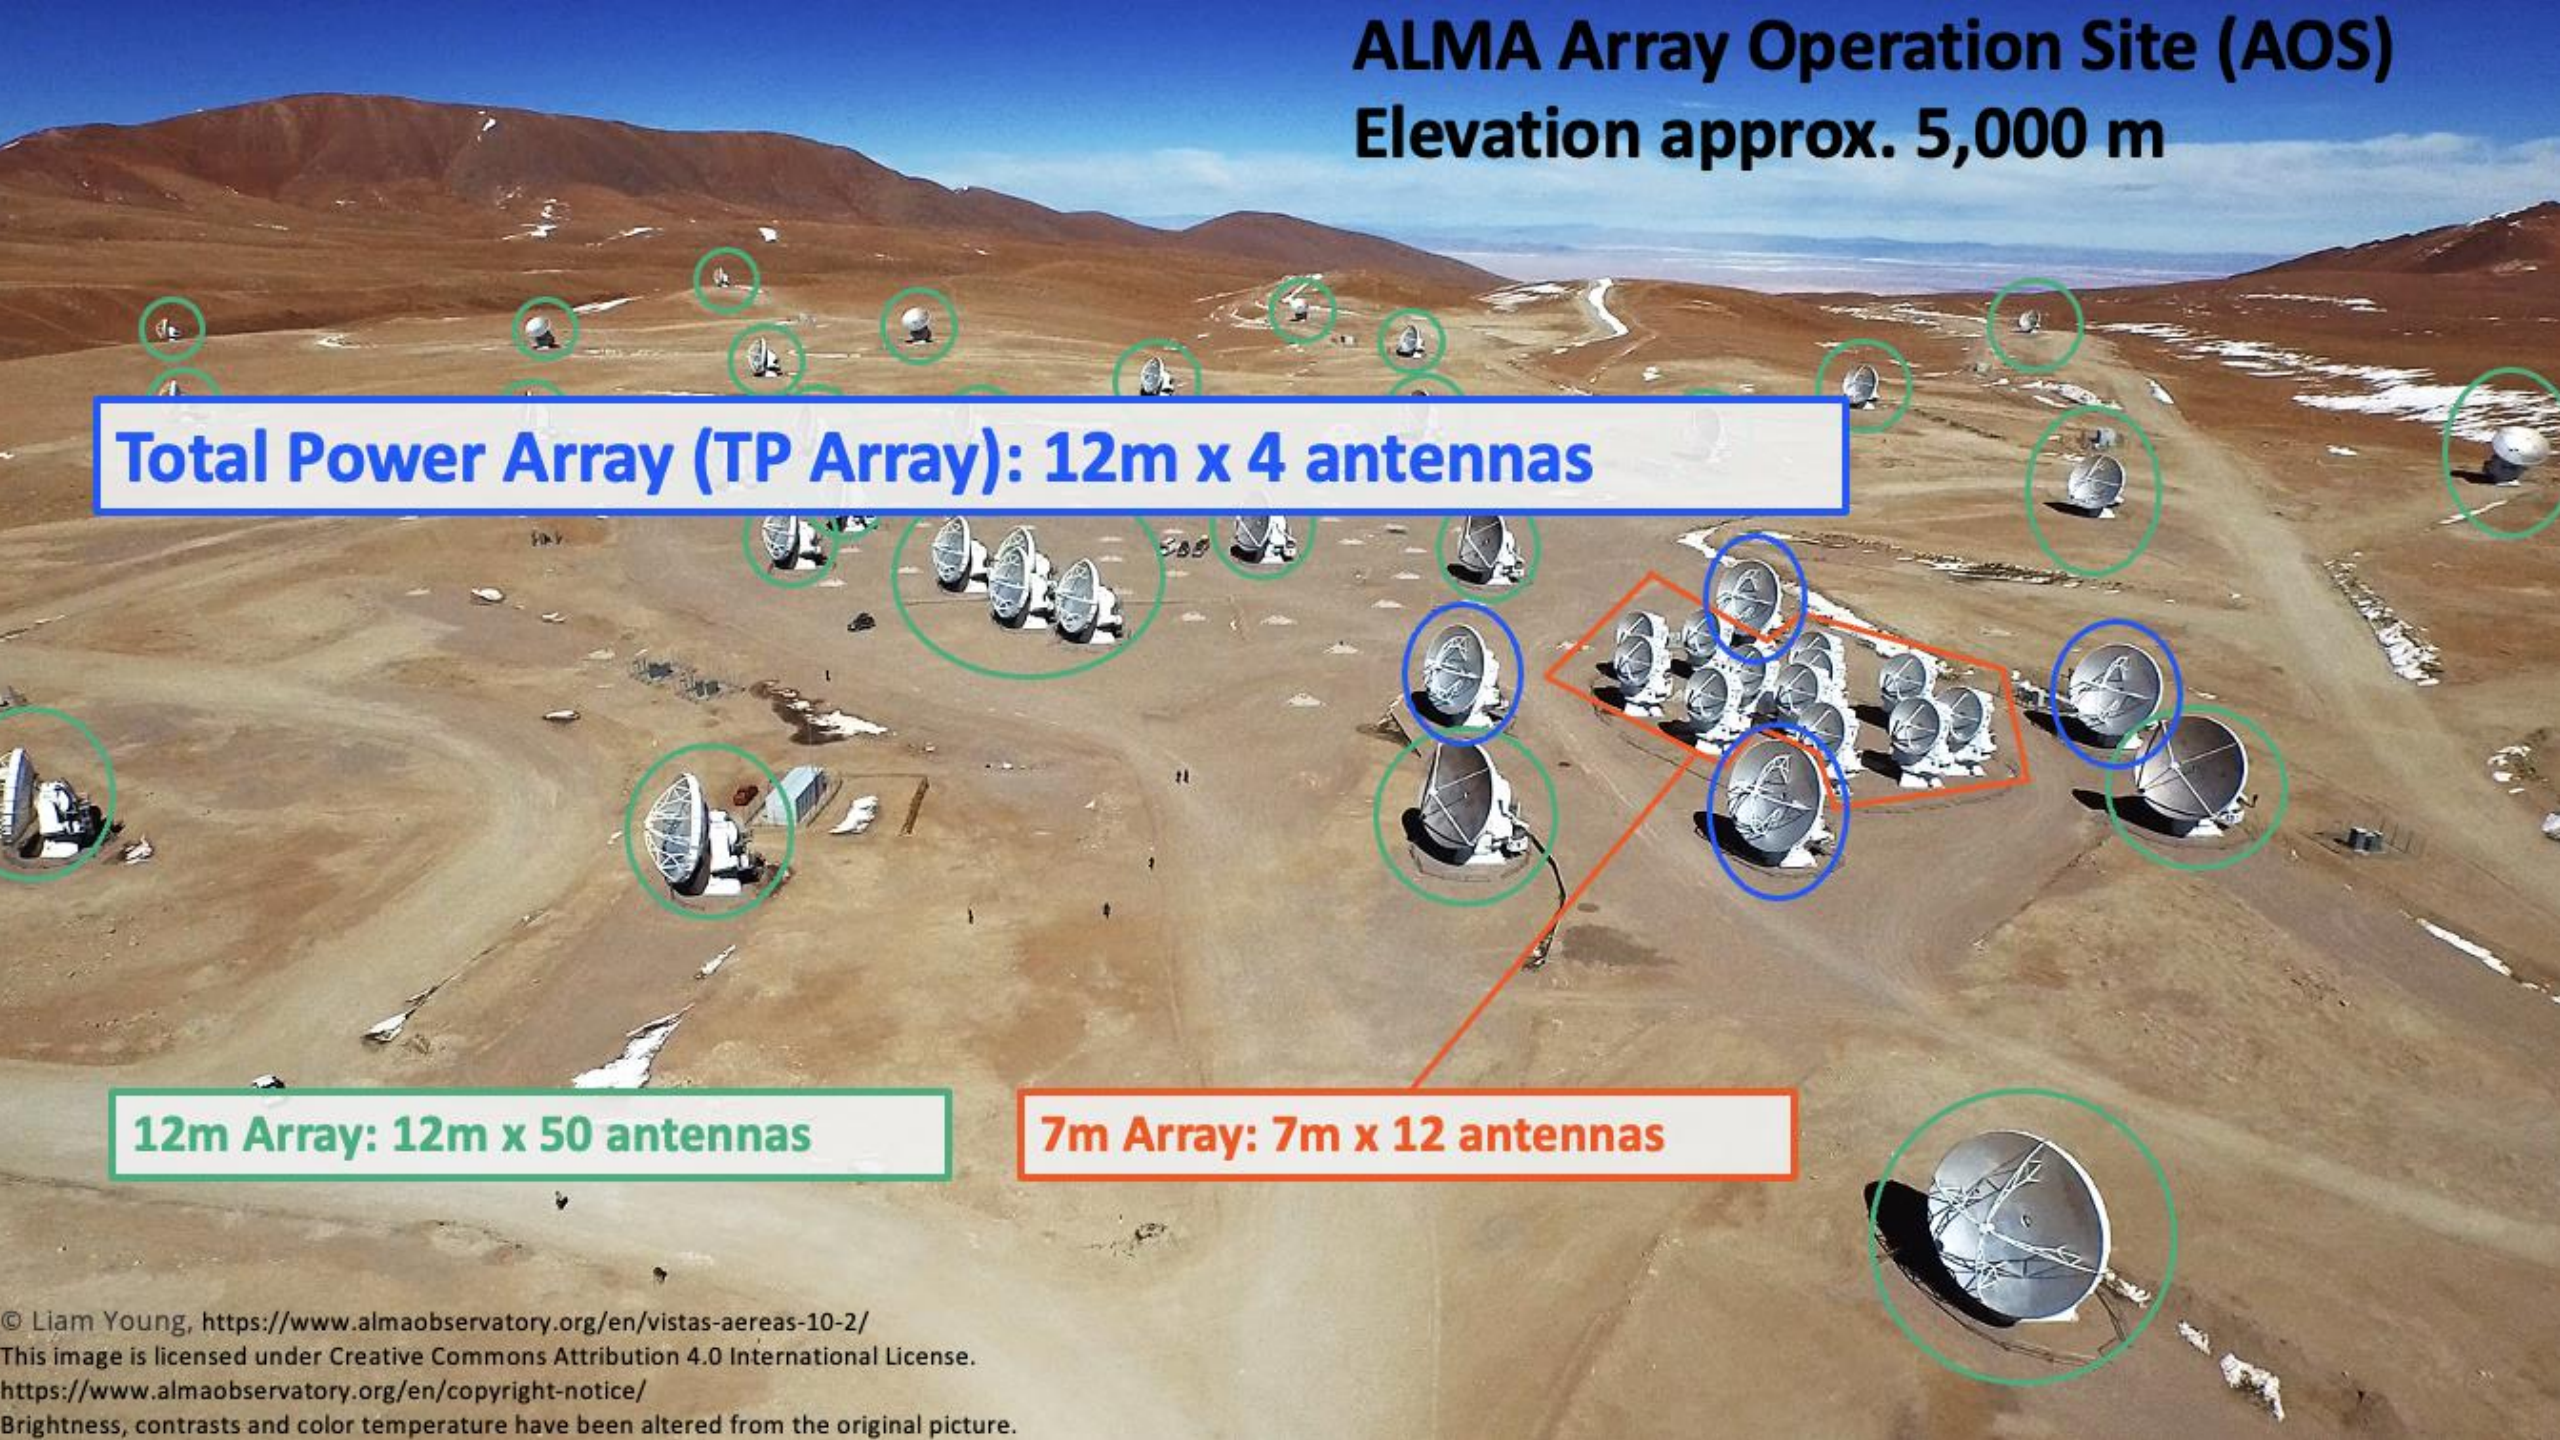

© Liam Young, https://www.almaobservatory.org/en/vistas-aereas-10-2/
This image is licensed under Creative Commons Attribution 4.0 International License.
https://www.almaobservatory.org/en/copyright-notice/
Brightness, contrasts and color temperature have been altered from the original picture.

ACA Spectrometer
Existing Spectrometer for Total Power Array

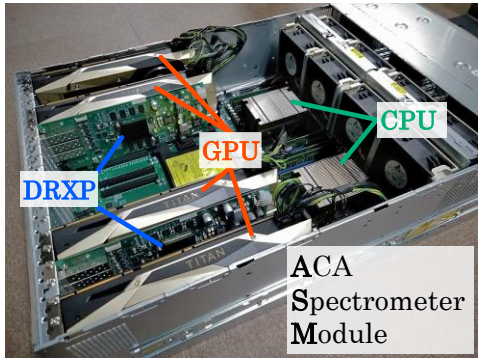

NVIDIA TITAN V
(16 GPUs in total)

Observing Time of Total Power Array per Cycle

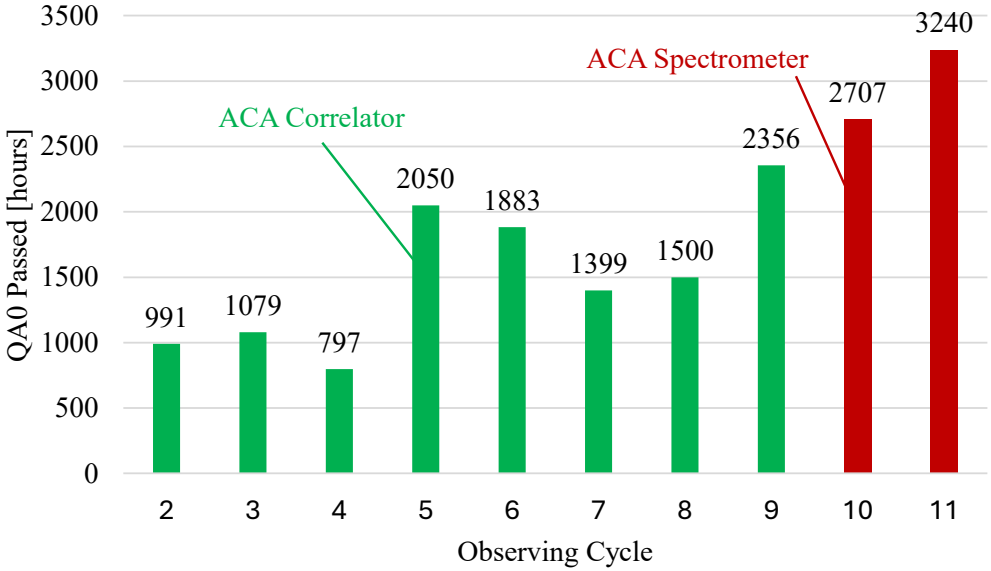

- Developed by KASI and NAOJ using GPU technologies.
- It was installed in AOS Technical Building (~5,000 m).
- It has been used in the science operation since Oct 2022 (Cycle 10) without significant issues.
- Will be replaced by Total Power GPU Spectrometer (TPGS) as a part of the Wideband Sensitivity Upgrade (WSU).

ACA Total Power Array Spectrometer Development Team Receives NAOJ Director General's Award (25/03)

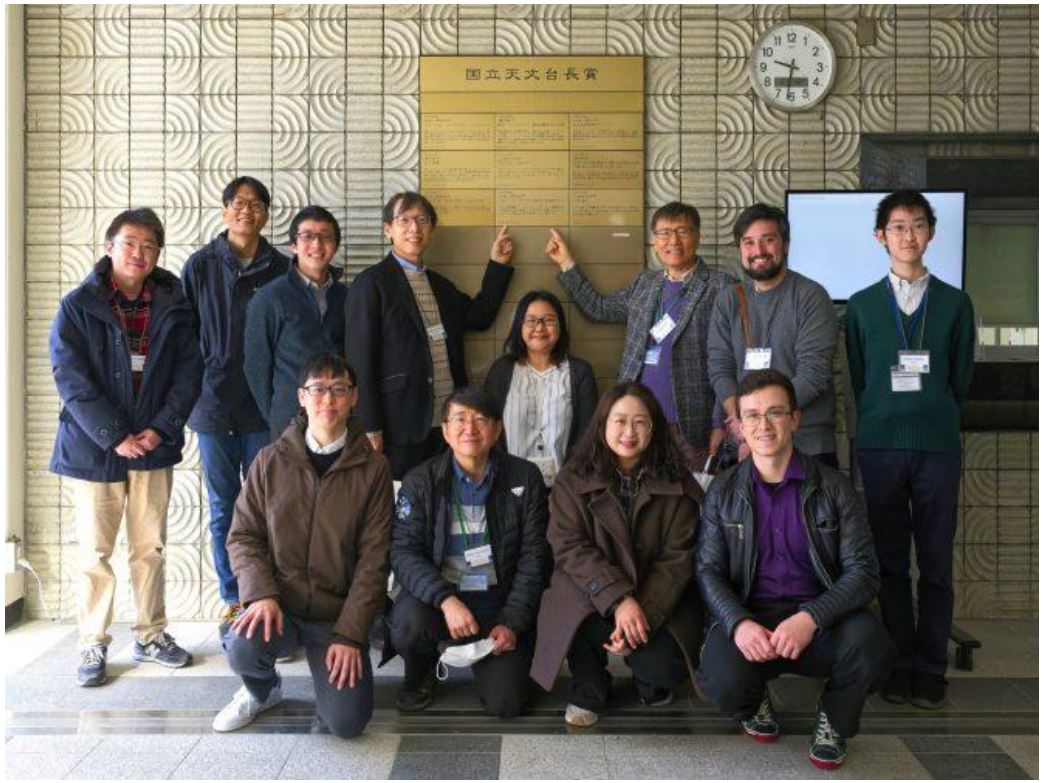

KASI: Jongsoo Kim, Jongsuk Hong, Won-Kee Park, Heesun Yoon
NAOJ: **Takashi Nakamoto**, Manabu Watanabe, Susumu Nakayama, Makoto Shizugami, Isaac Goicovich, Shinya Ikebe, Camilo Saldias

Wideband Sensitivity Upgrade (WSU)

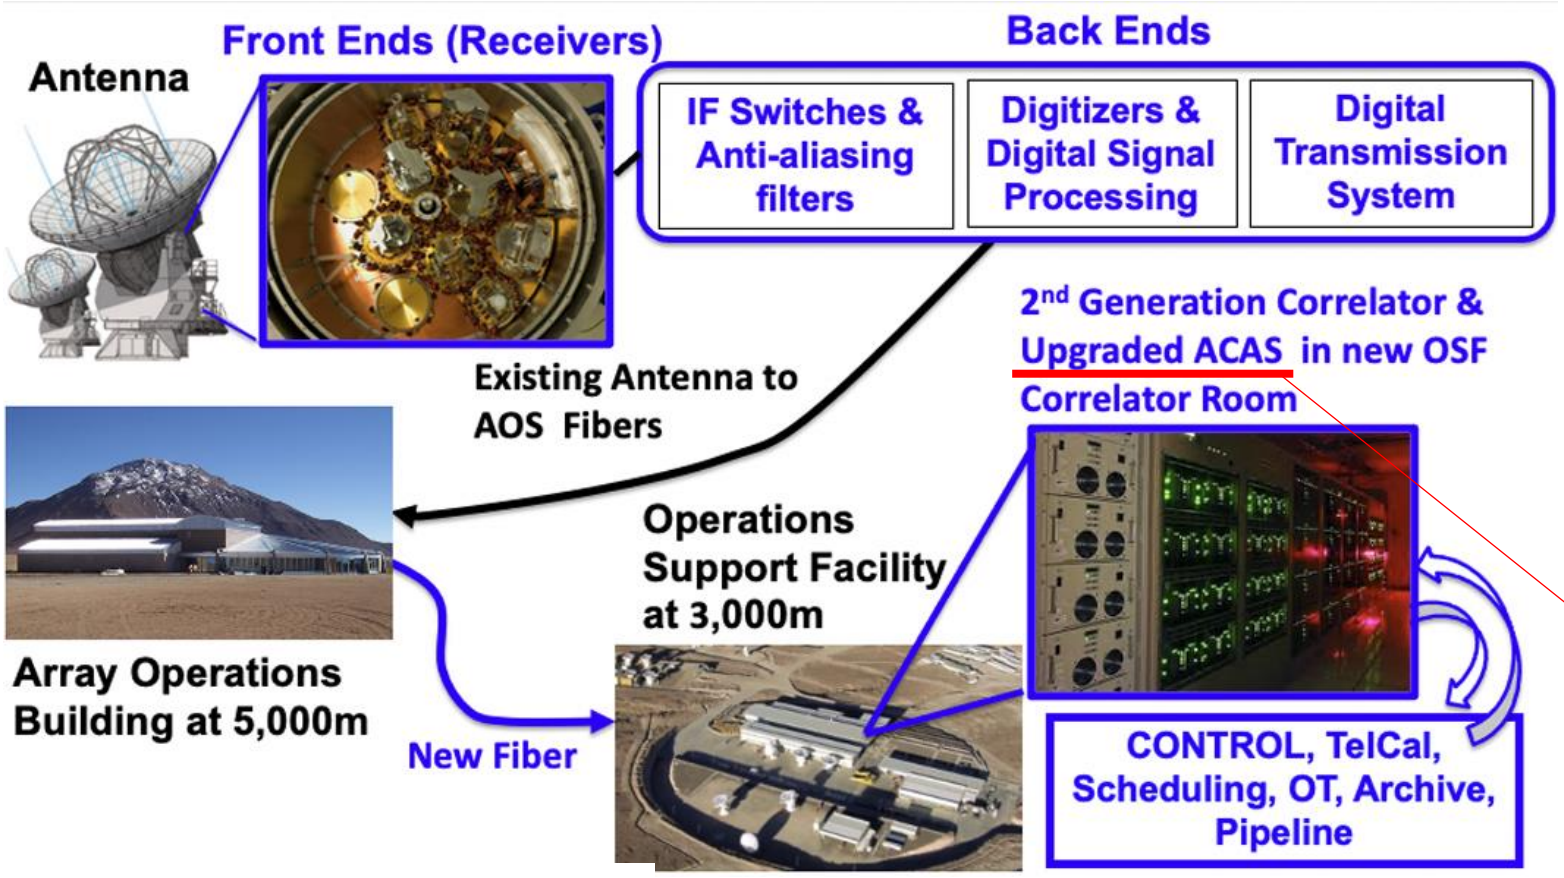

New components or components that will be upgraded are in blue.

Many major components in ALMA will be replaced or upgraded.

TPGS will be installed in OSF, not AOS.

Reference:
John Carpenter et al.,
The ALMA2030 Wideband Sensitivity Upgrade
ALMA Memo 621 (https://library.nrao.edu/public/memos/alma/main/memo621.pdf)

TPGS Key Technical Specifications and Technology Choices

Items	ACA Spectrometer	TPGS
Sampling rate and number of inputs per antenna	4 Gsps x 4 BB x 2 pol	40 Gsps x 2 SB x 2 pol
Number of bits per sample	3 bits	6 bits
Total input data rate from four antennas	384 Gbps	3.84 Tbps
Maximum correlator bandwidth (CBW) per pol	8 GHz	32 GHz
Channel width at max CBW (full pol)	976 kHz	13.5 kHz
Finest channel width (full pol)	61.04 kHz	1.7 kHz

Fully utilize up-to-date GPUs and 400 GbE technologies to achieve the high input data rate, the high bandwidth, and the fine channel resolutions.

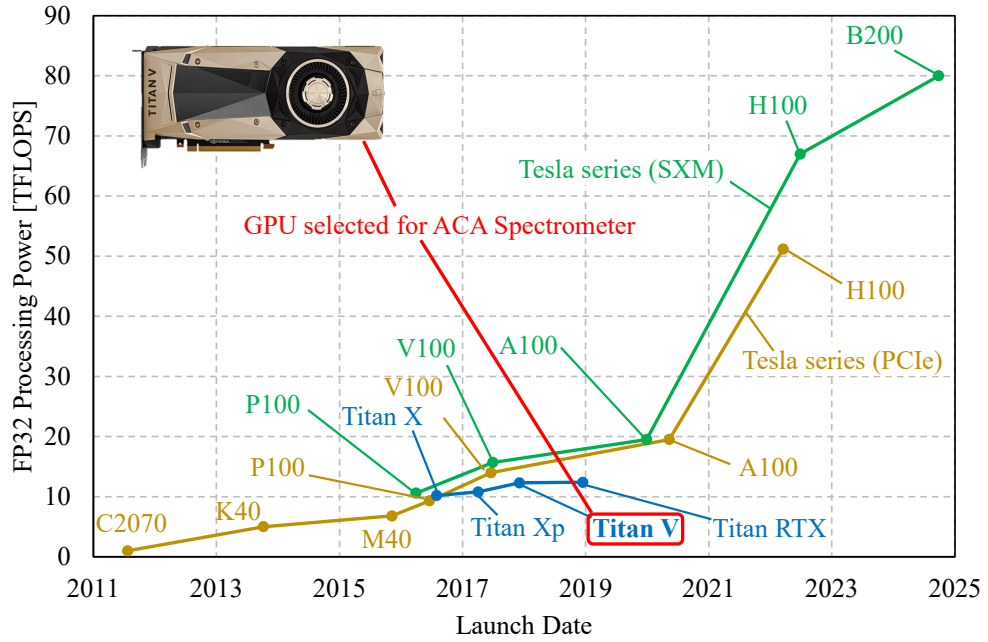

TPGS System Architecture

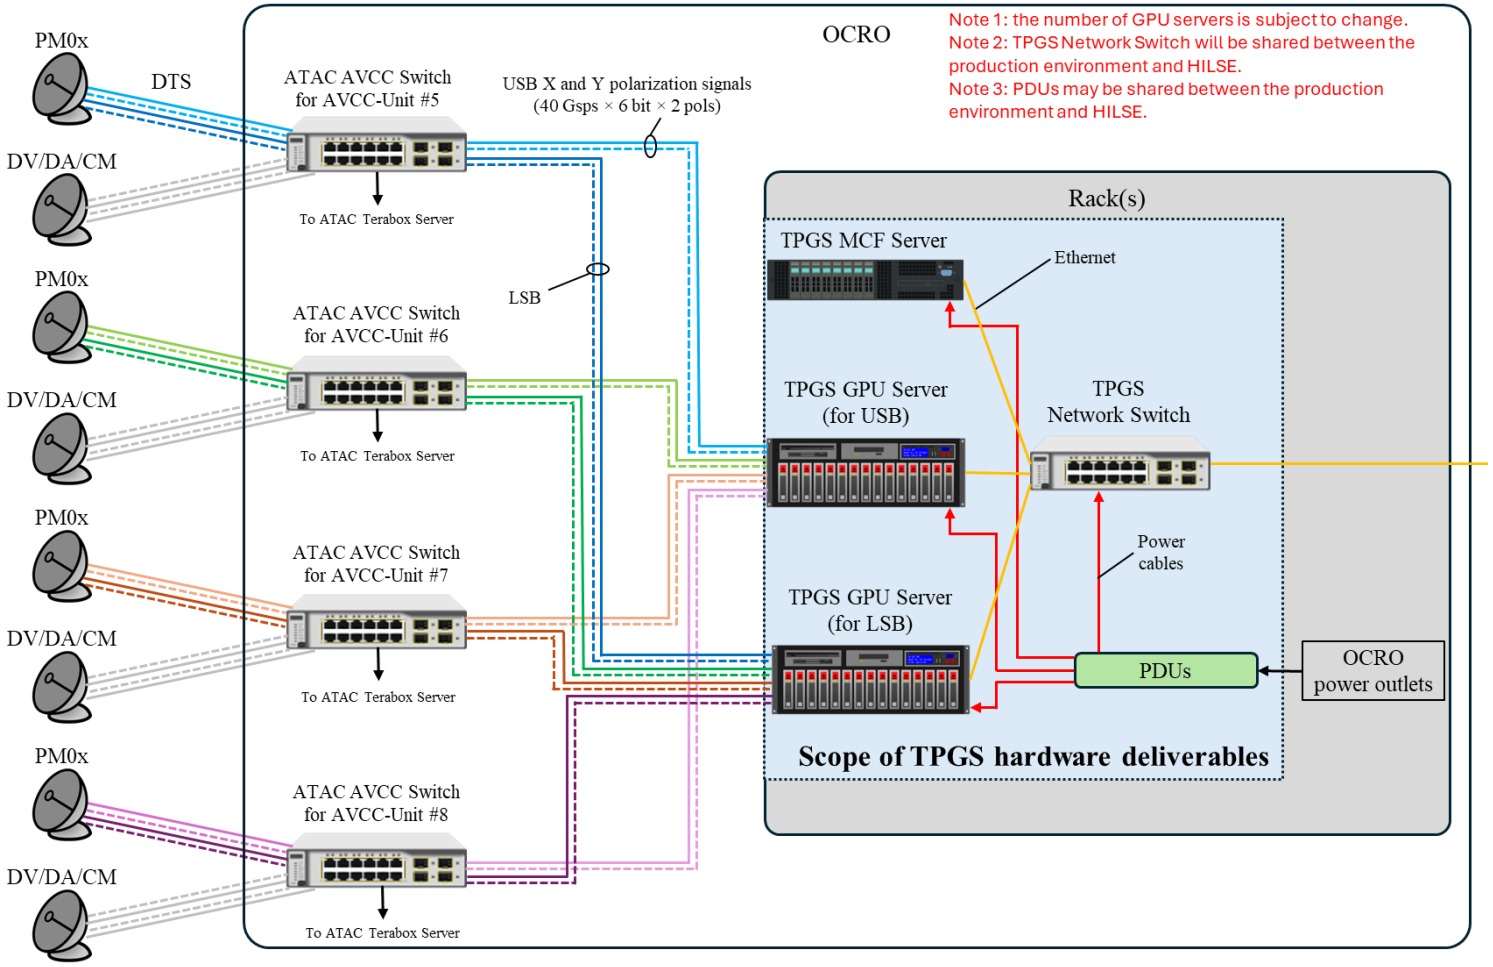

Two GPU servers, one for USB and the other for LSB.

One GPU server receives two polarization signals from four antennas, which needs eight 400 GbE ports.

Note: This is our current baseline design. The design is subject to change.

Server Block Diagram with 8 data streams

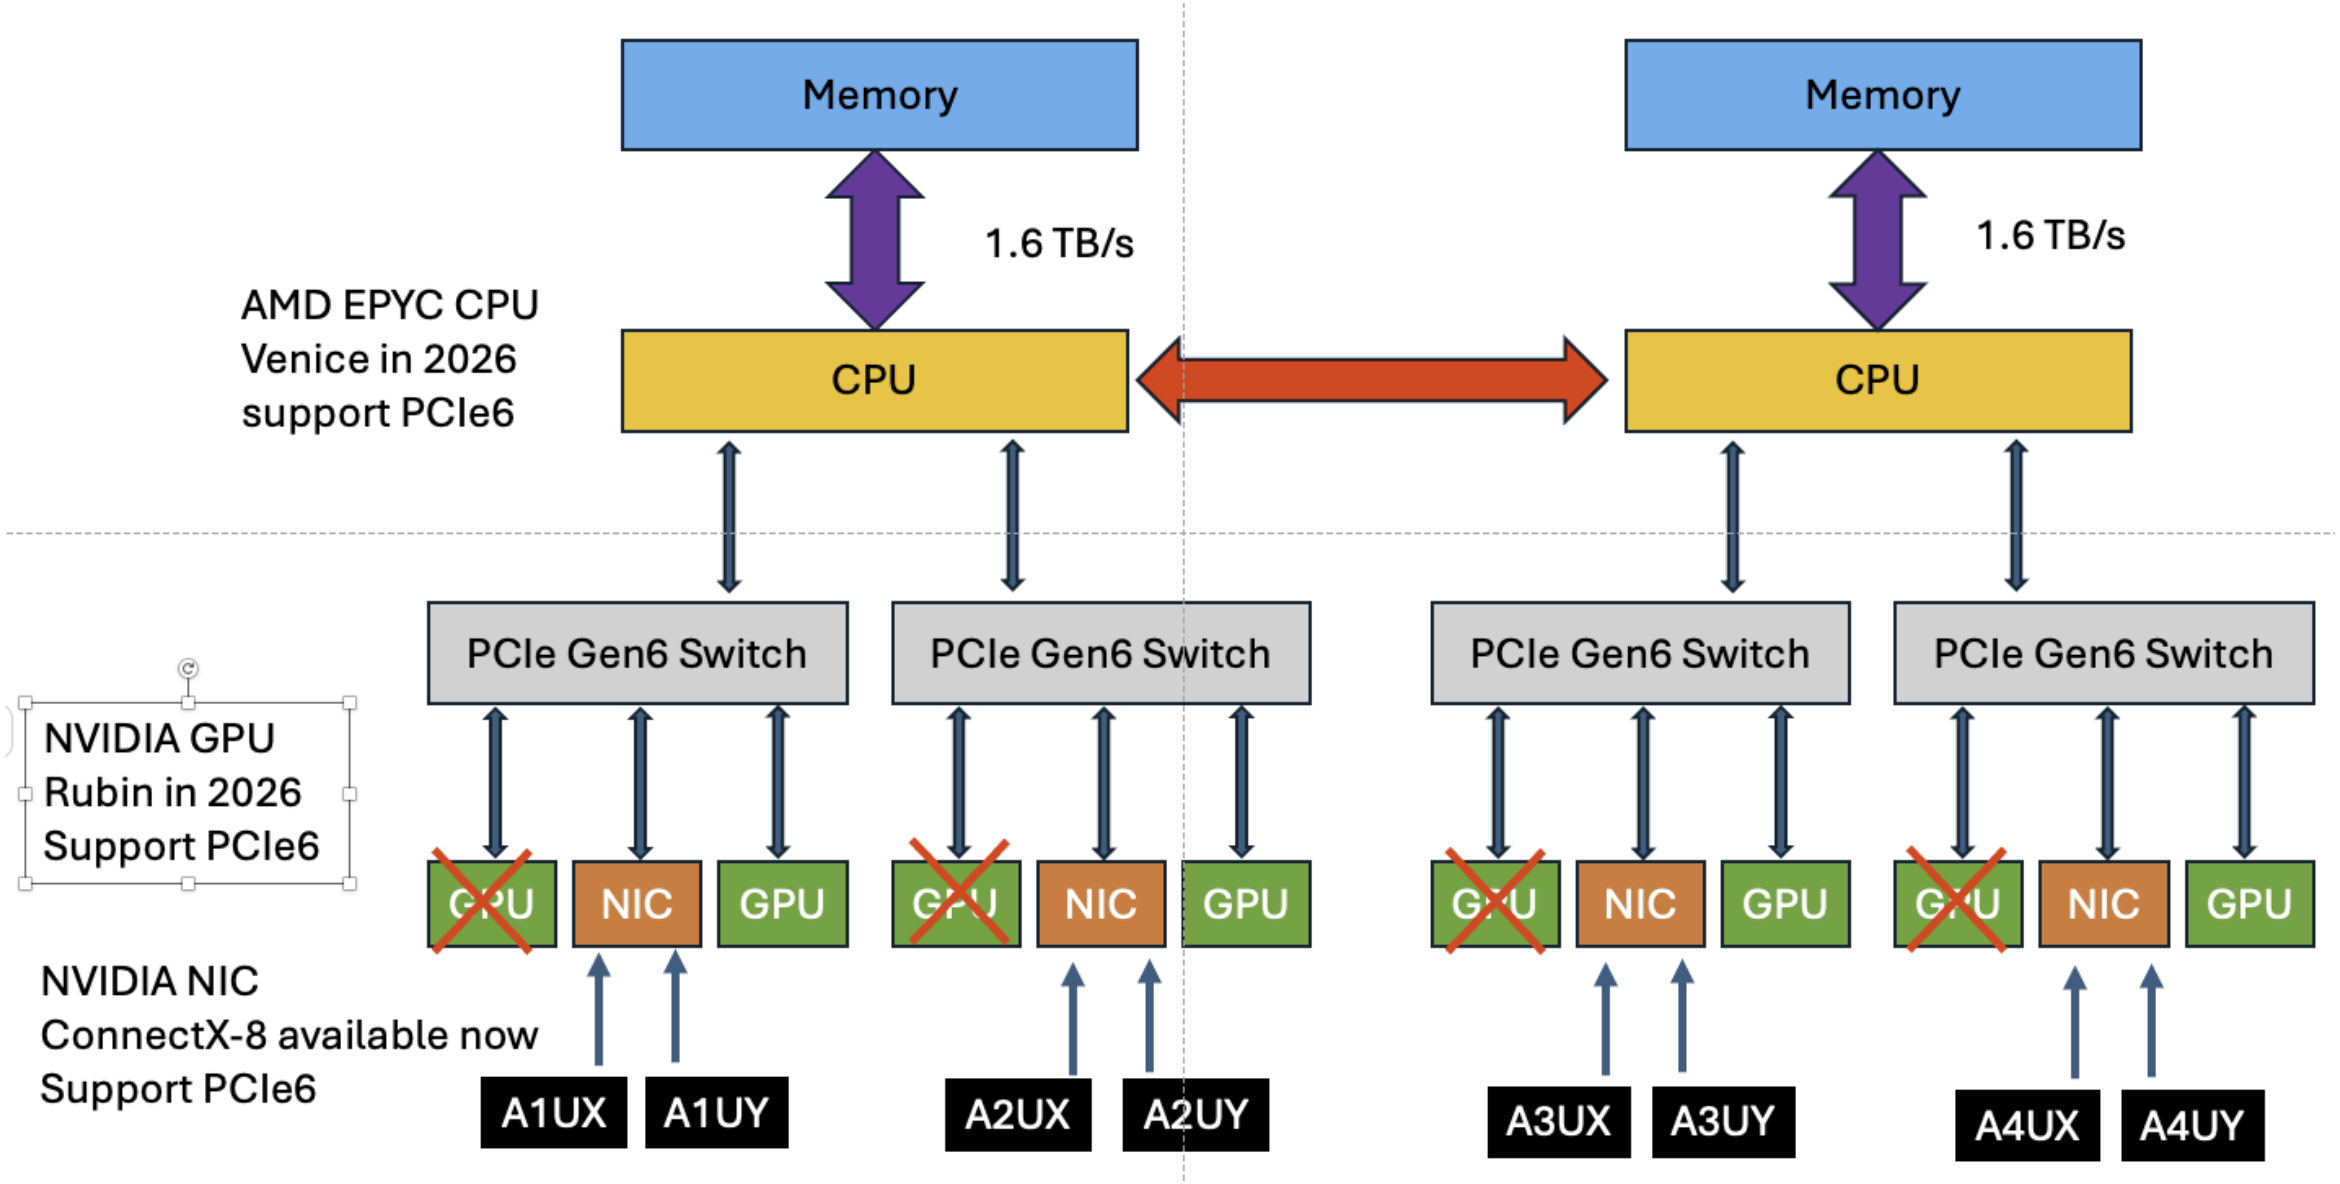

TPGS Software Architecture

TPGS-CU (Control Unit)
TPGS-PU (Processing Unit)
TPGS-MU (Monitoring Unit)

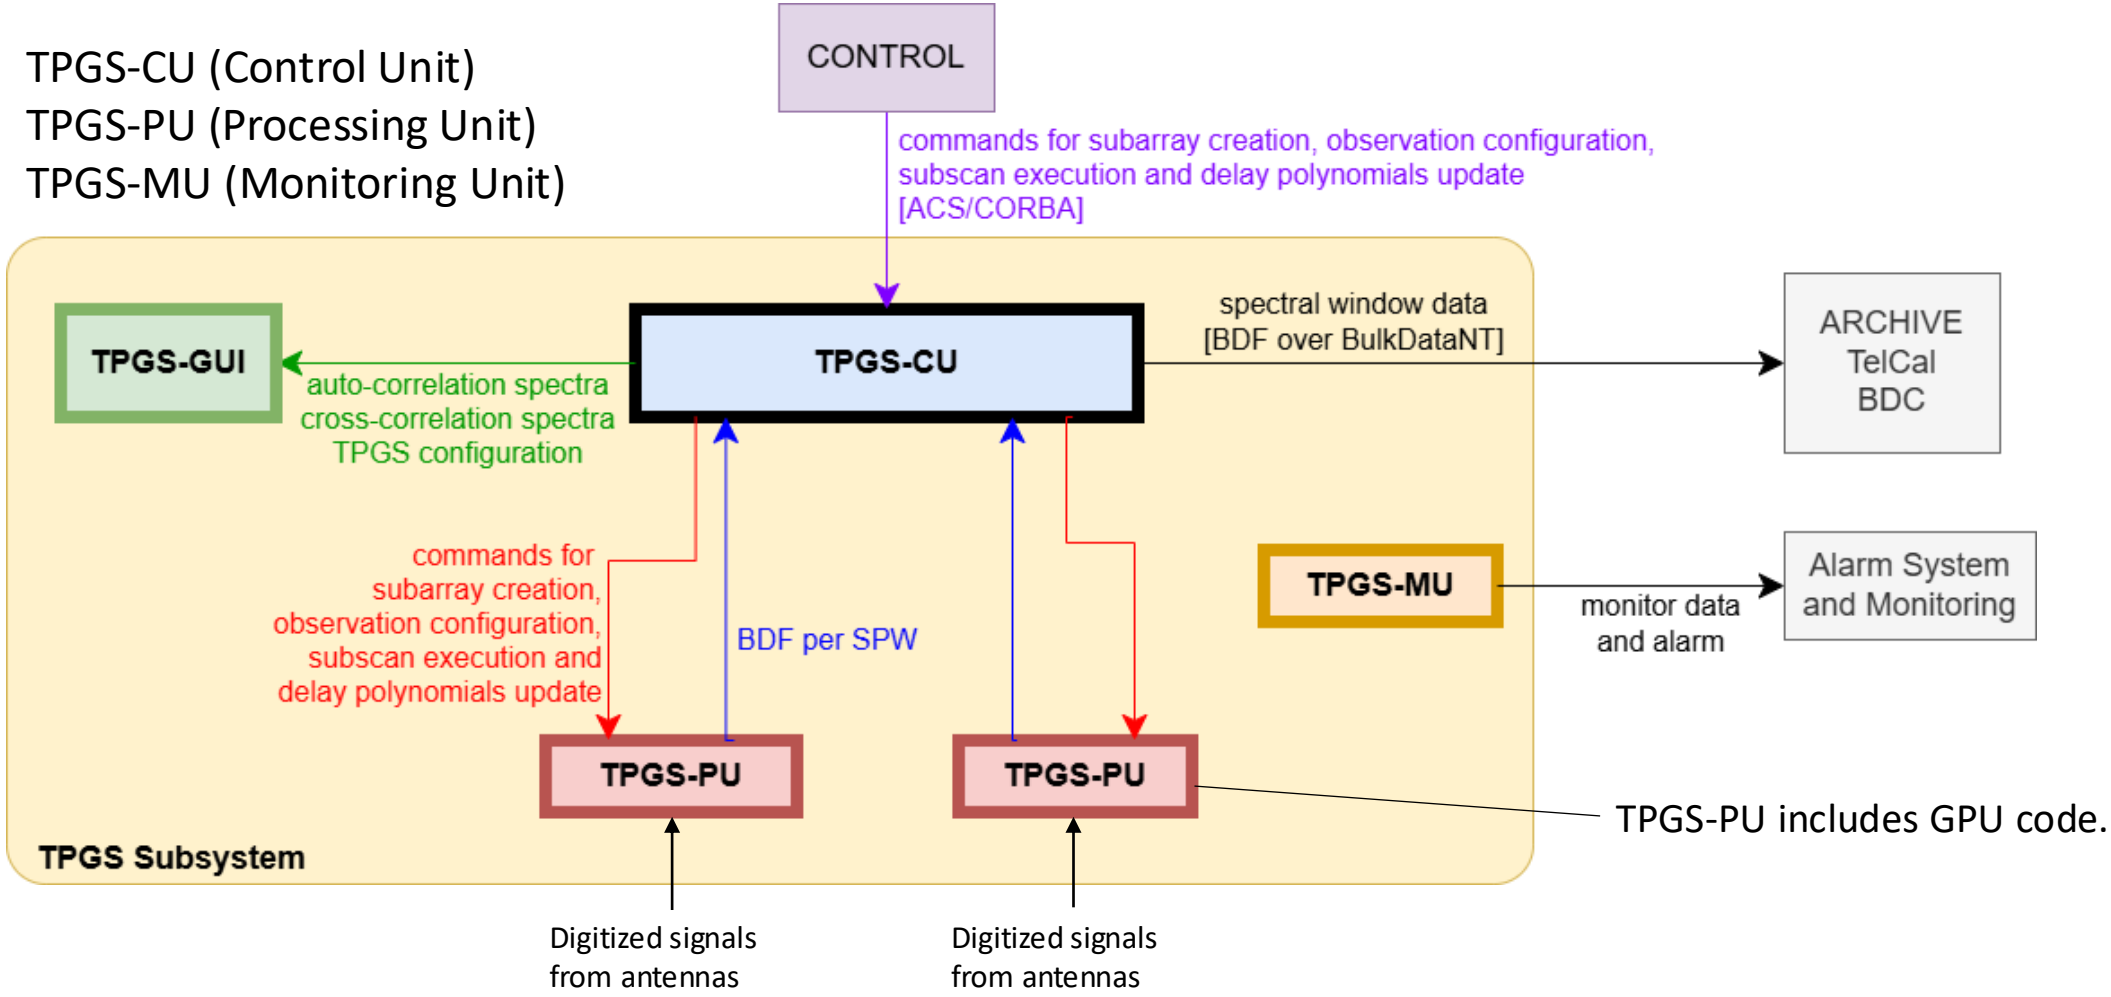

Note: Design is subject to change.

Dual Pipelines: Auto & Cross-Correlation

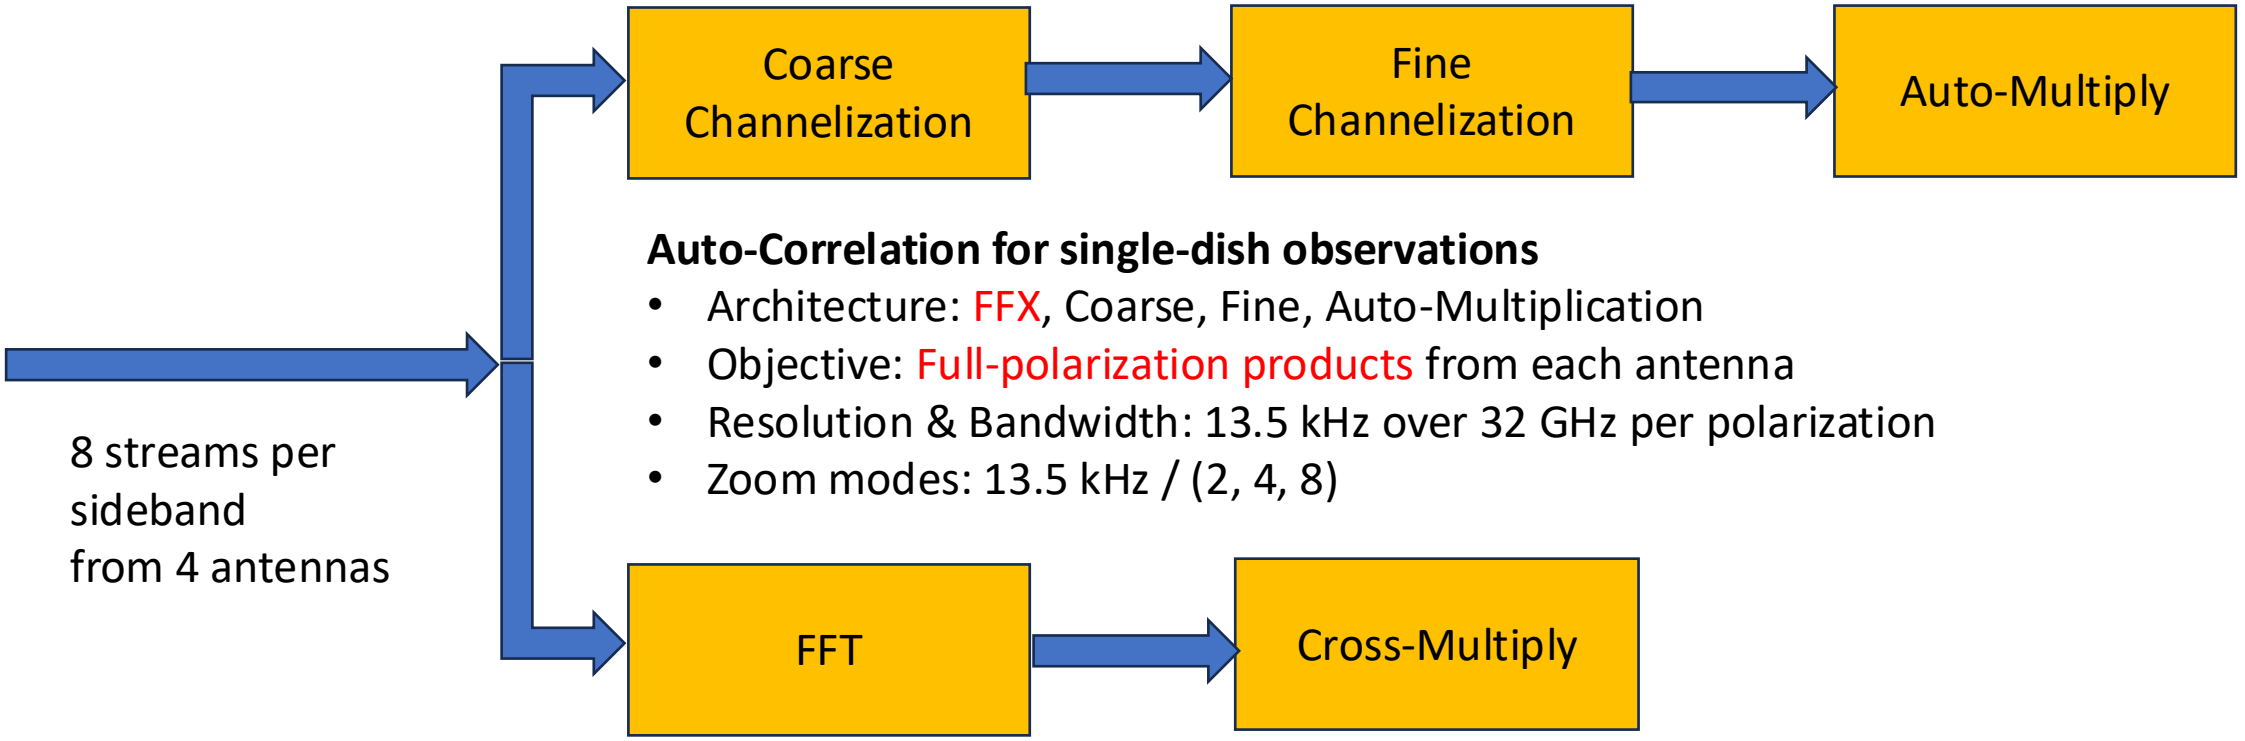

Auto-Correlation for single-dish observations

- Architecture: FFX, Coarse, Fine, Auto-Multiplication
- Objective: Full-polarization products from each antenna
- Resolution & Bandwidth: 13.5 kHz over 32 GHz per polarization
- Zoom modes: 13.5 kHz / (2, 4, 8)

Cross-Correlation for pointing and focus calibration

- Architecture: FX, FFT and Cross-Multiplication
- Objective: Dual-polarization products from four antennas
- Resolution: 20 MHz over a 2 GHz SPW; 4 SPWs per sideband
- No quantization and online WVR corrections

Polyphase Filter Bank and a Prototype Filter

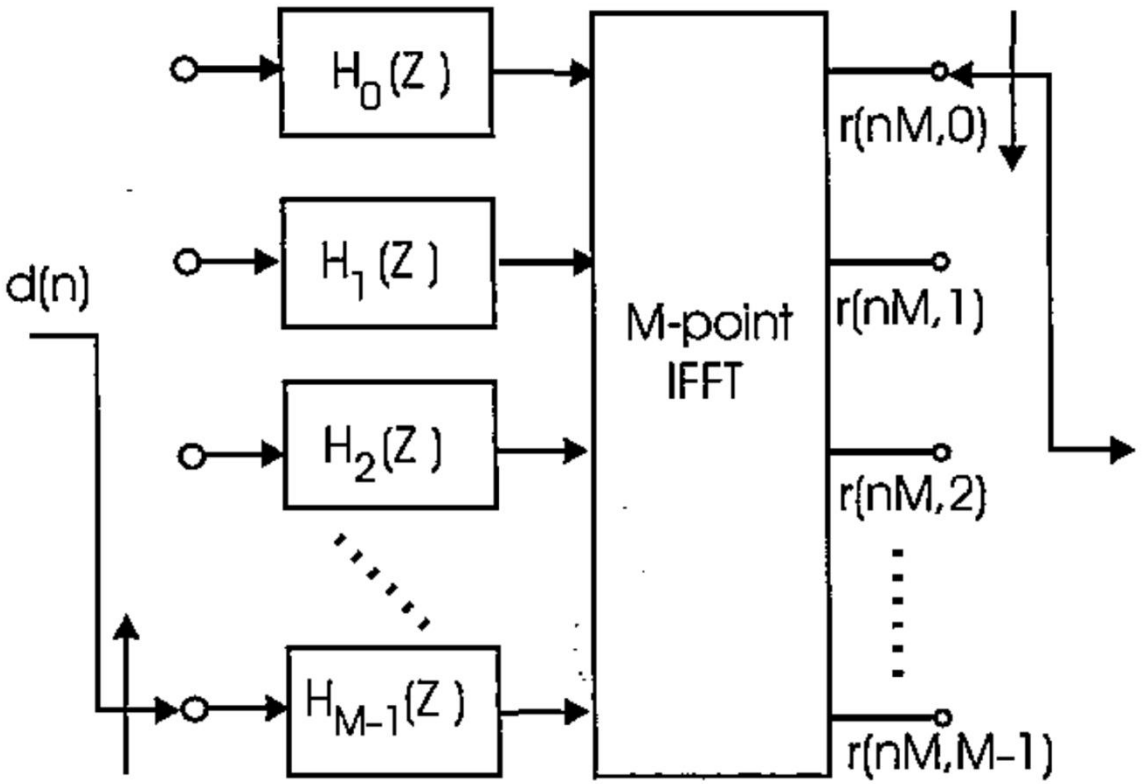

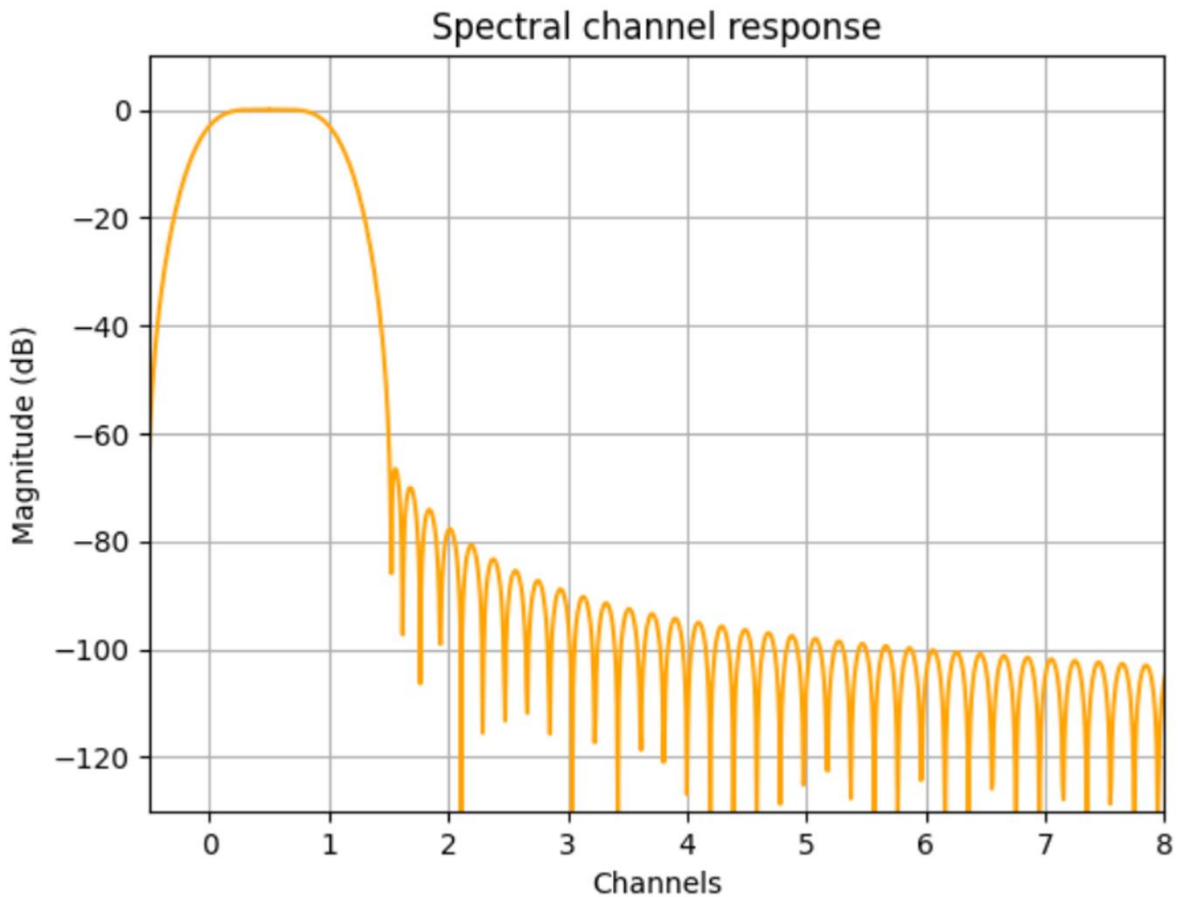

Jongsuk Hong's Poster

Coarse OSPFB (Dual Polarization)

40 GS/s ADC → 6-bit quantize → Polyphase FIR → 2048-pt FFT → Phase correction → Hermitian separation

ADC Input

40 GS/s × 2 pol (X, Y)

640 M samples (16 ms frame)

Signal Generation

Gaussian noise (σ = 8) + CW tones

820 tones/pol, amp ∈ [0.02, 0.30]

6-bit Quantize + Pack

clip [−32, +31] → 4 samples / uint32

s_0[23:18] s_1[17:12] s_2[11:6] s_3[5:0]

Histogram

64-bin (fused in kernel)

Gaussian σ verification

Data Rates

ADC: 40 GS/s × 6 bit = 30 GB/s / pol

Packed: 2 × 153 MB (16 ms)

Throughput: 60 GB/s aggregate

Coarse OSPFB Kernel

M = 2048 D = 1448 T = 6 (oversampled)

Polyphase FIR → 2048-pt C2C FFT → Phase Correction

EPT = 4, FFTsPerBlock = 2, cuFFTDx (SM 1200)

Dual-pol trick: X → Re, Y → Im

One FFT handles both polarizations

Phase $\varphi_k(m) = -2\pi k \cdot m \cdot D/M$

Phasor renorm every 1024 iters

Hermitian Separation

$$X = \tfrac{1}{2}\left(Z_k + Z^*_{-k} \right)$$

$$Y = \tfrac{1}{2j}\left(Z^*_{-k} - Z_k \right)$$

Shared memory exchange: 2 × M × cuComplex (32 KB)

Bank-conflict free shared memory layout

Lane 0 → smem[0..M-1], Lane 1 → smem[M..2M-1]

16 GHz Sub-band Output

820 of 1025 bins $k \in [103, 922]$

`Frame-major: d_xpol[ m × 820 + k ]`

440,320 frames × 820 ch × cuComplex / pol

Coalesced global memory writes

Output size: 2 × 2.72 GB (X and Y polarizations)

→ p.2

GPU Kernels — Page 1

K0 `gen_gaussian`

K1 `coarse_ospfb_dualpol`

Tiled Transpose (Dual Polarization)

Frame-major layout → 32×32 tiled transpose → Subband-major layout

Frame-Major Buffer (Coarse Output)

`d_xpol[ m × 820 + k ]`

Rows = 440,320 frames Cols = 820 subbands

Size: 440,320 × 820 × 8 B
= 2.72 GB per polarization

Access: row-major reads
Thread block reads one tile row

Transpose Summary

Input:	Frame-major [440K × 820]
Output:	Subband-major [820 × 440K]
Tile:	32 × 32 with +1 padding
Bandwidth:	~5.44 GB transferred (2 pol)
Kernel:	K2 transpose_kernel (×2 pol)

Tiled Transpose Kernel

TILE = 32 × 32 shared memory: cuComplex tile[32][33]

Read tile row-major (coalesced) → Store to smem → Write column-major (coalesced)

Padding (+1 column) eliminates shared memory bank conflicts

→ p.3

Grid: [820/32] × [440320/32] = 26 × 13760 Block: (32, 32) = 1024 threads
Both read and write are fully coalesced

Executed twice: once for X-pol, once for Y-pol
In-place not possible — separate src and dst buffers

Subband-Major Buffer (Fine Input)

`d_xpol_t[ sb × N_frames + m ]`

Rows = 820 subbands Cols = 440,320 frames

1 CUDA block = 1 subband → coalesced temporal reads

Fine CSPFB + Power Spectra (Dual Polarization)

Subband-major input → Polyphase FIR → 2048-pt FFT → Passband select → XX / YY / XY power

Subband-Major Input

`d_xpol_t[ sb × 440320 + m ]`

1 CUDA block = 1 subband, threadIdx.y = pol (0=X, 1=Y)

Resolution Summary

ADC bandwidth:	20 GHz (Nyquist)
Coarse Δf:	19.53 MHz (40G / 2048)
Fine Δf:	13.49 kHz (27.6M / 2048)
Selected BW:	16 GHz (820 subbands)
Total fine ch:	820 × 1448 = 1,187,360

Fine CSPFB Kernel

M = 2048   T = 5   (critically sampled, D = M)

Polyphase FIR → 2048-pt C2C FFT → Auto/Cross power accumulation

EPT = 4, FFTsPerBlock = 2 (pol 0=X, pol 1=Y), cuFFTDx (SM 1200)

Passband Selection + fftshift

Keep D/M = 1448 / 2048 = 70.7% of each sub-FFT

FFT ch [0, 724) → out [724, 1448)   (positive freq)

FFT ch [1324, 2048) → out [0, 724)   (negative freq)

820 blocks × (512, 2) threads

215 frames integrated per subband

Cross-power via smem exchange

$X_k \cdot Y^*_k$ accumulated per fine channel

Brick-wall subband boundaries

Guard bands discarded cleanly

Δf = 13.49 kHz (final resolution)

Total BW = 16 GHz

Power Spectra Output

820 subbands × 1448 channels = 1,187,360 fine channels

XX (Auto-Power)

$\Sigma\ |X_k|^2$ per channel (float)

YY (Auto-Power)

$\Sigma\ |Y_k|^2$ per channel (float)

XY (Cross-Power)

$\Sigma\ X_k \cdot Y^*_k$ per channel (cuComplex)

GPU Specification Comparison

	RTX PRO 6000 Blackwell Server	GH200 Grace Hopper
Architecture	Blackwell (SM 12.0)	Hopper (SM 9.0)
SMs	188	132
CUDA Cores	24,064	16,896
FP32 Cores / SM	128	128
GPU Clock (locked)	2,430 MHz	1,980 MHz
FP32 Peak	**117 TFLOPS**	**67 TFLOPS**
Memory Type	GDDR7	HBM3e
Memory Size	96 GB	144 GB
Memory Bandwidth	**1.6 TB/s**	**4.9 TB/s**
Ridge Point (AI)	73.1 FLOP/B	13.7 FLOP/B
Bus Width	512-bit	6144-bit
TDP	600 W	700 W
Form Factor	PCIe	Superchip (module)

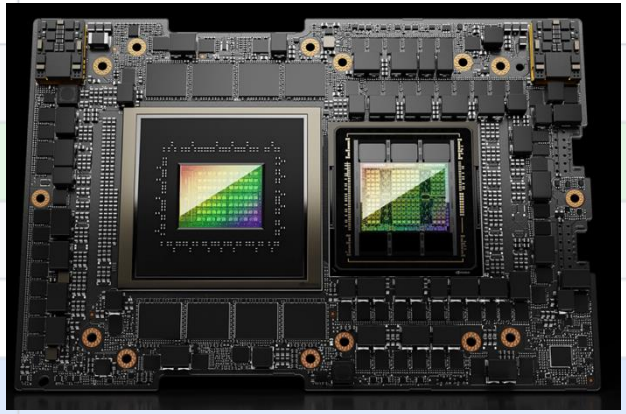

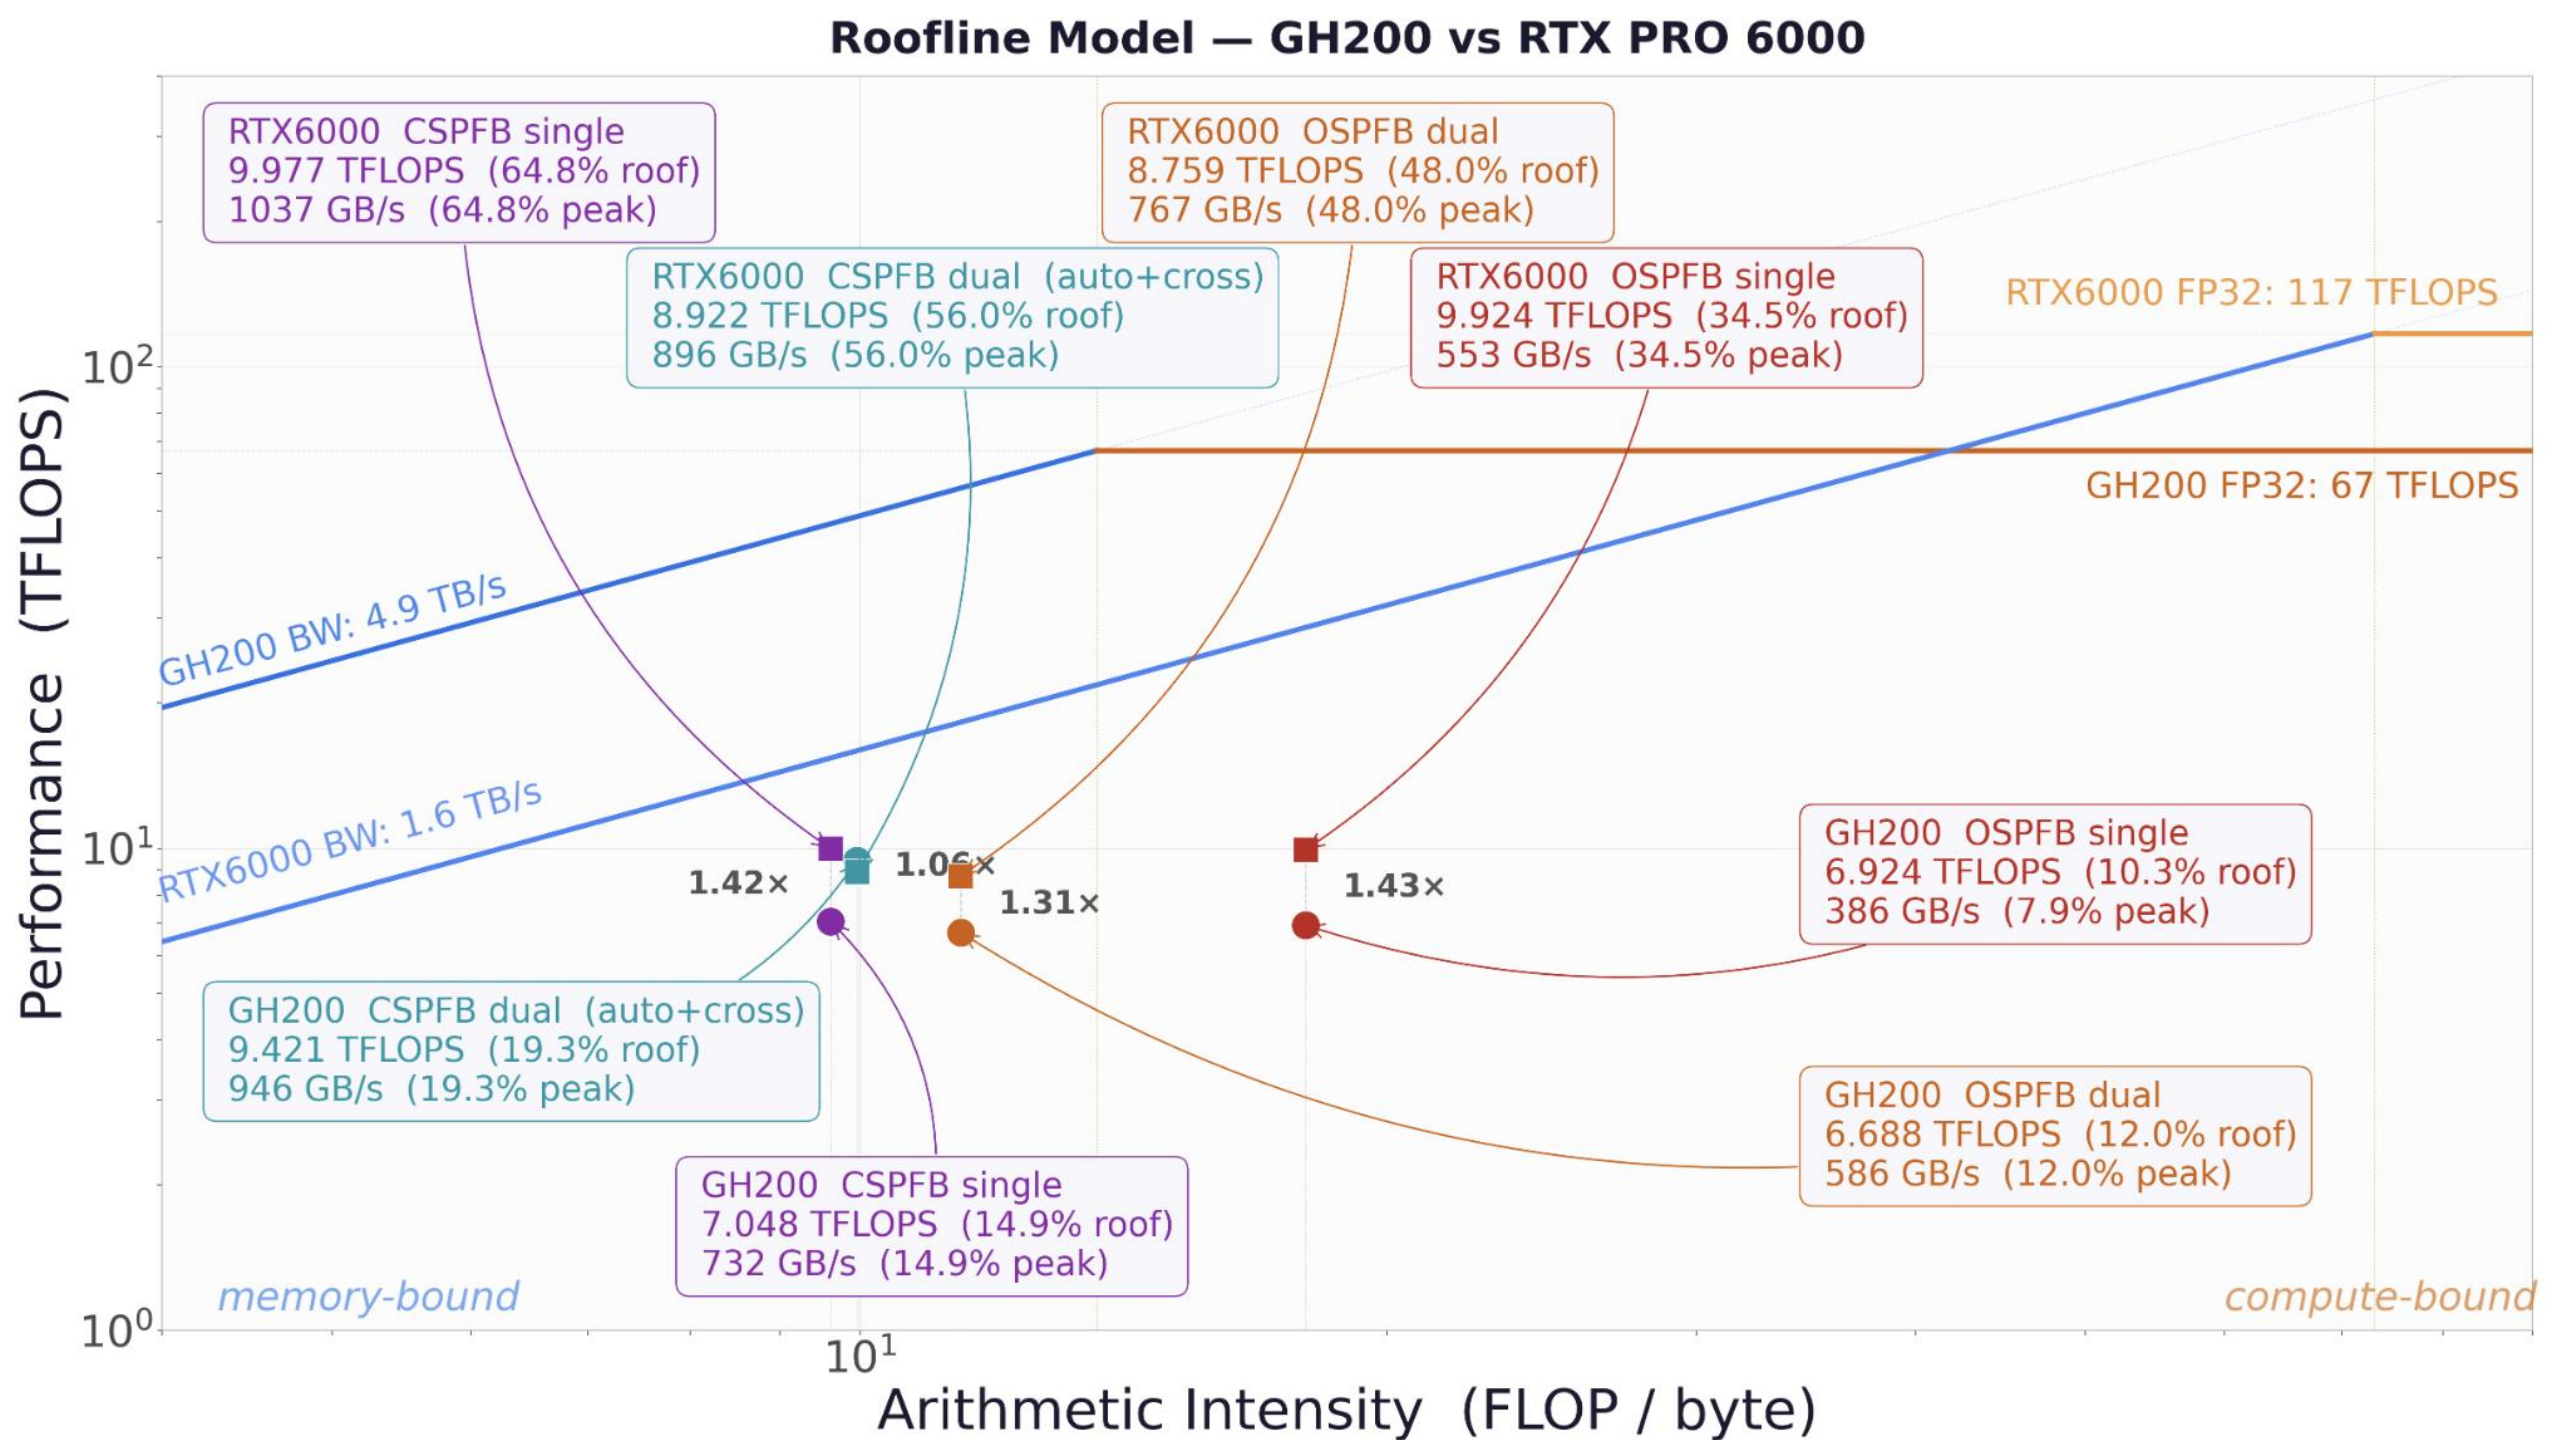
Roofline Model — GH200 vs RTX PRO 6000
RTX6000 CSPFB single
9.977 TFLOPS (64.8% roof)
1037 GB/s (64.8% peak)
RTX6000 CSPFB dual (auto+cross)
8.922 TFLOPS (56.0% roof)
896 GB/s (56.0% peak)
RTX6000 OSPFB dual
8.759 TFLOPS (48.0% roof)
767 GB/s (48.0% peak)
RTX6000 OSPFB single
9.924 TFLOPS (34.5% roof)
553 GB/s (34.5% peak)
RTX6000 FP32: 117 TFLOPS
GH200 FP32: 67 TFLOPS
GH200 BW: 4.9 TB/s
RTX6000 BW: 1.6 TB/s
1.42×
1.06×
1.31×
1.43×
GH200 OSPFB single
6.924 TFLOPS (10.3% roof)
386 GB/s (7.9% peak)
GH200 CSPFB dual (auto+cross)
9.421 TFLOPS (19.3% roof)
946 GB/s (19.3% peak)
GH200 OSPFB dual
6.688 TFLOPS (12.0% roof)
586 GB/s (12.0% peak)
GH200 CSPFB single
7.048 TFLOPS (14.9% roof)
732 GB/s (14.9% peak)
memory-bound
compute-bound
Performance (TFLOPS)
Arithmetic Intensity (FLOP / byte)
10^2
10^1
10^0
10^1

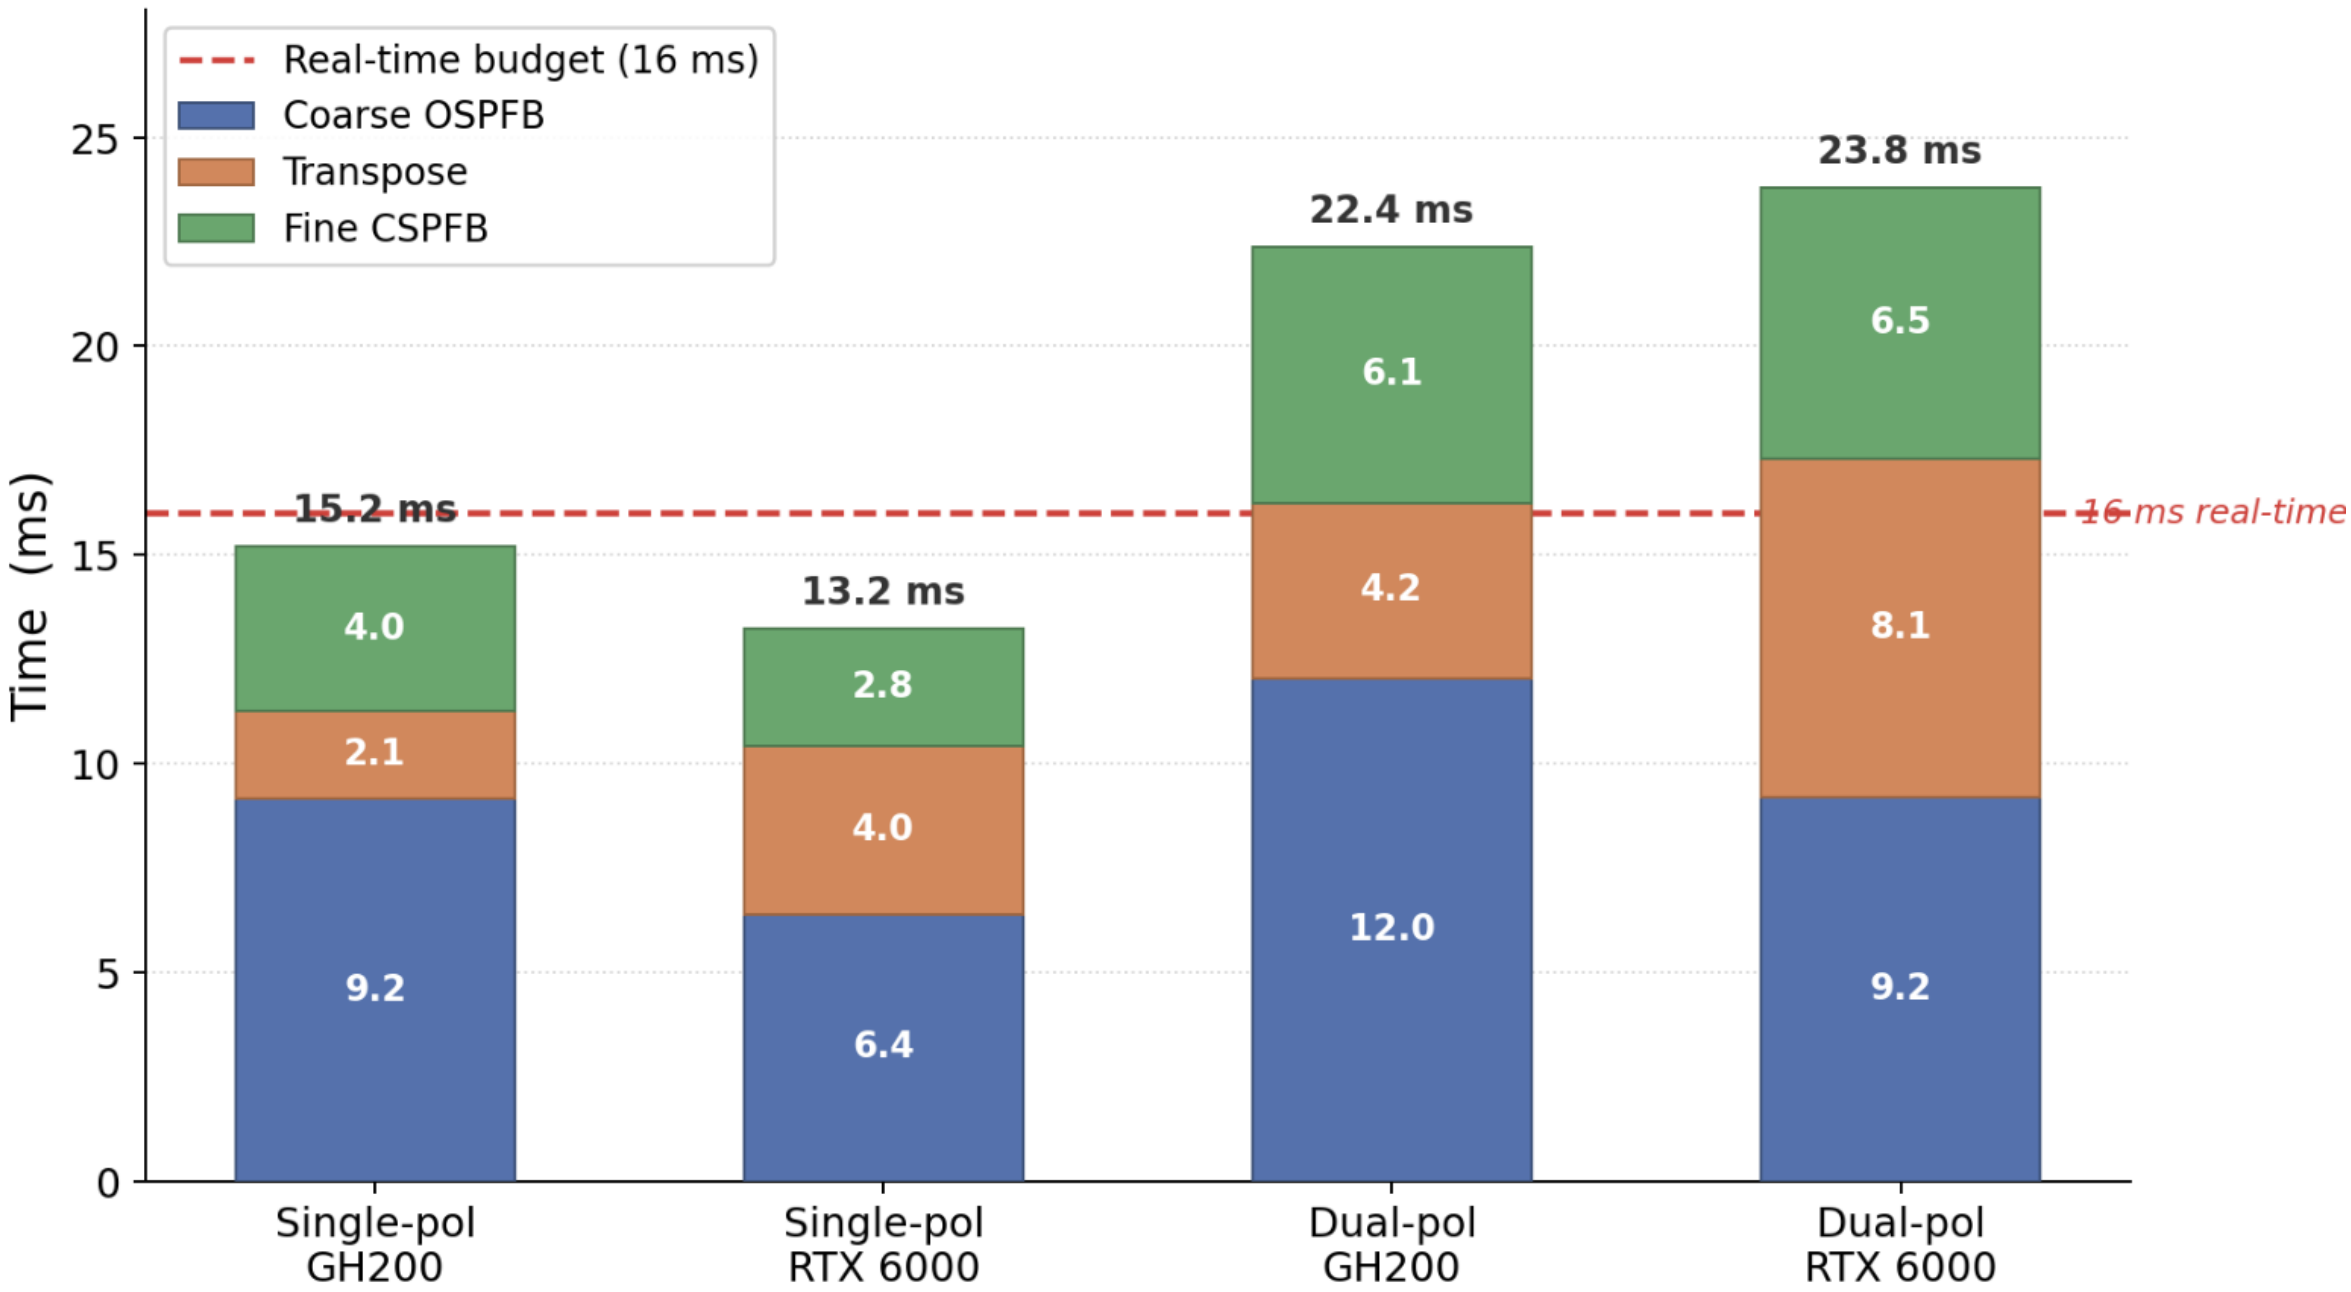
Single-Dish Pipeline — Kernel Timing Breakdown
Real-time budget (16 ms)
Coarse OSPFB
Transpose
Fine CSPFB
Time (ms)
0
5
10
15
20
25
15.2 ms
4.0
2.1
9.2
13.2 ms
2.8
4.0
6.4
22.4 ms
6.1
4.2
12.0
23.8 ms
6.5
8.1
9.2
16 ms real-time
Single-pol GH200
Single-pol RTX 6000
Dual-pol GH200
Dual-pol RTX 6000

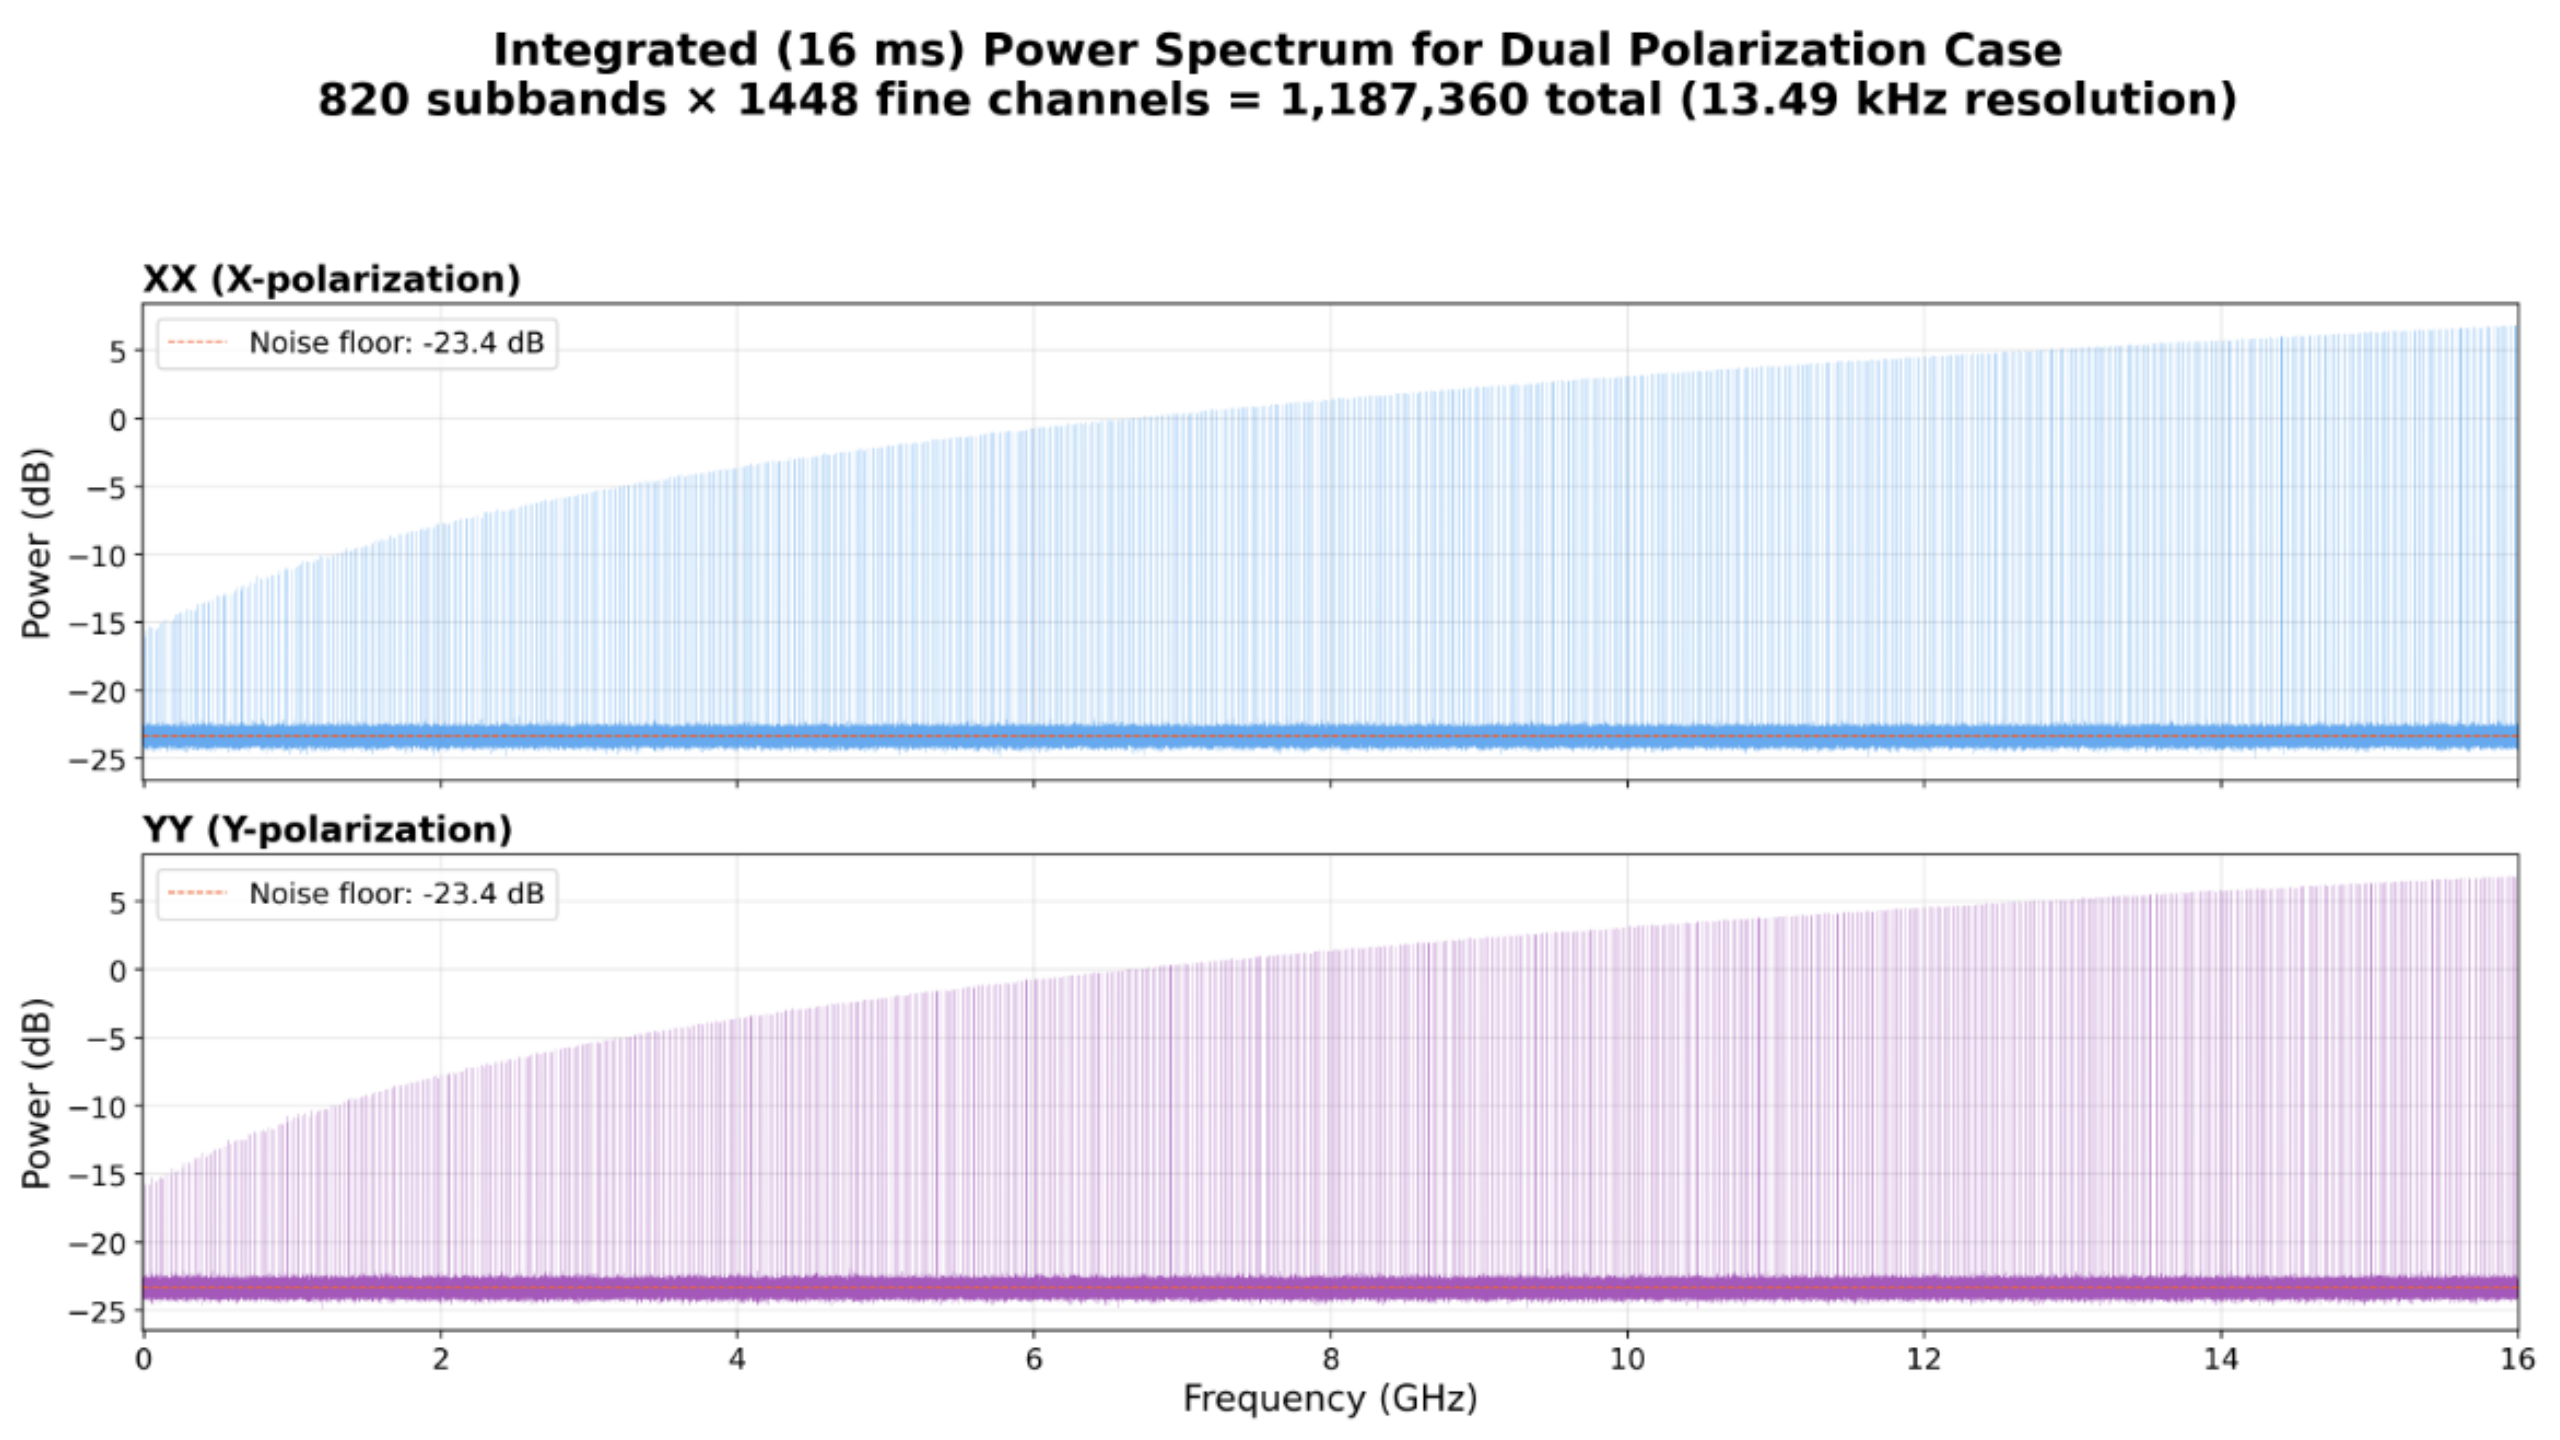
Integrated (16 ms) Power Spectrum for Dual Polarization Case
820 subbands × 1448 fine channels = 1,187,360 total (13.49 kHz resolution)
XX (X-polarization)
Noise floor: -23.4 dB
Power (dB)
5
0
−5
−10
−15
−20
−25
YY (Y-polarization)
Noise floor: -23.4 dB
Power (dB)
5
0
−5
−10
−15
−20
−25
0
2
4
6
8
10
12
14
16
Frequency (GHz)

Subband Spectra (XX top, YY bottom)

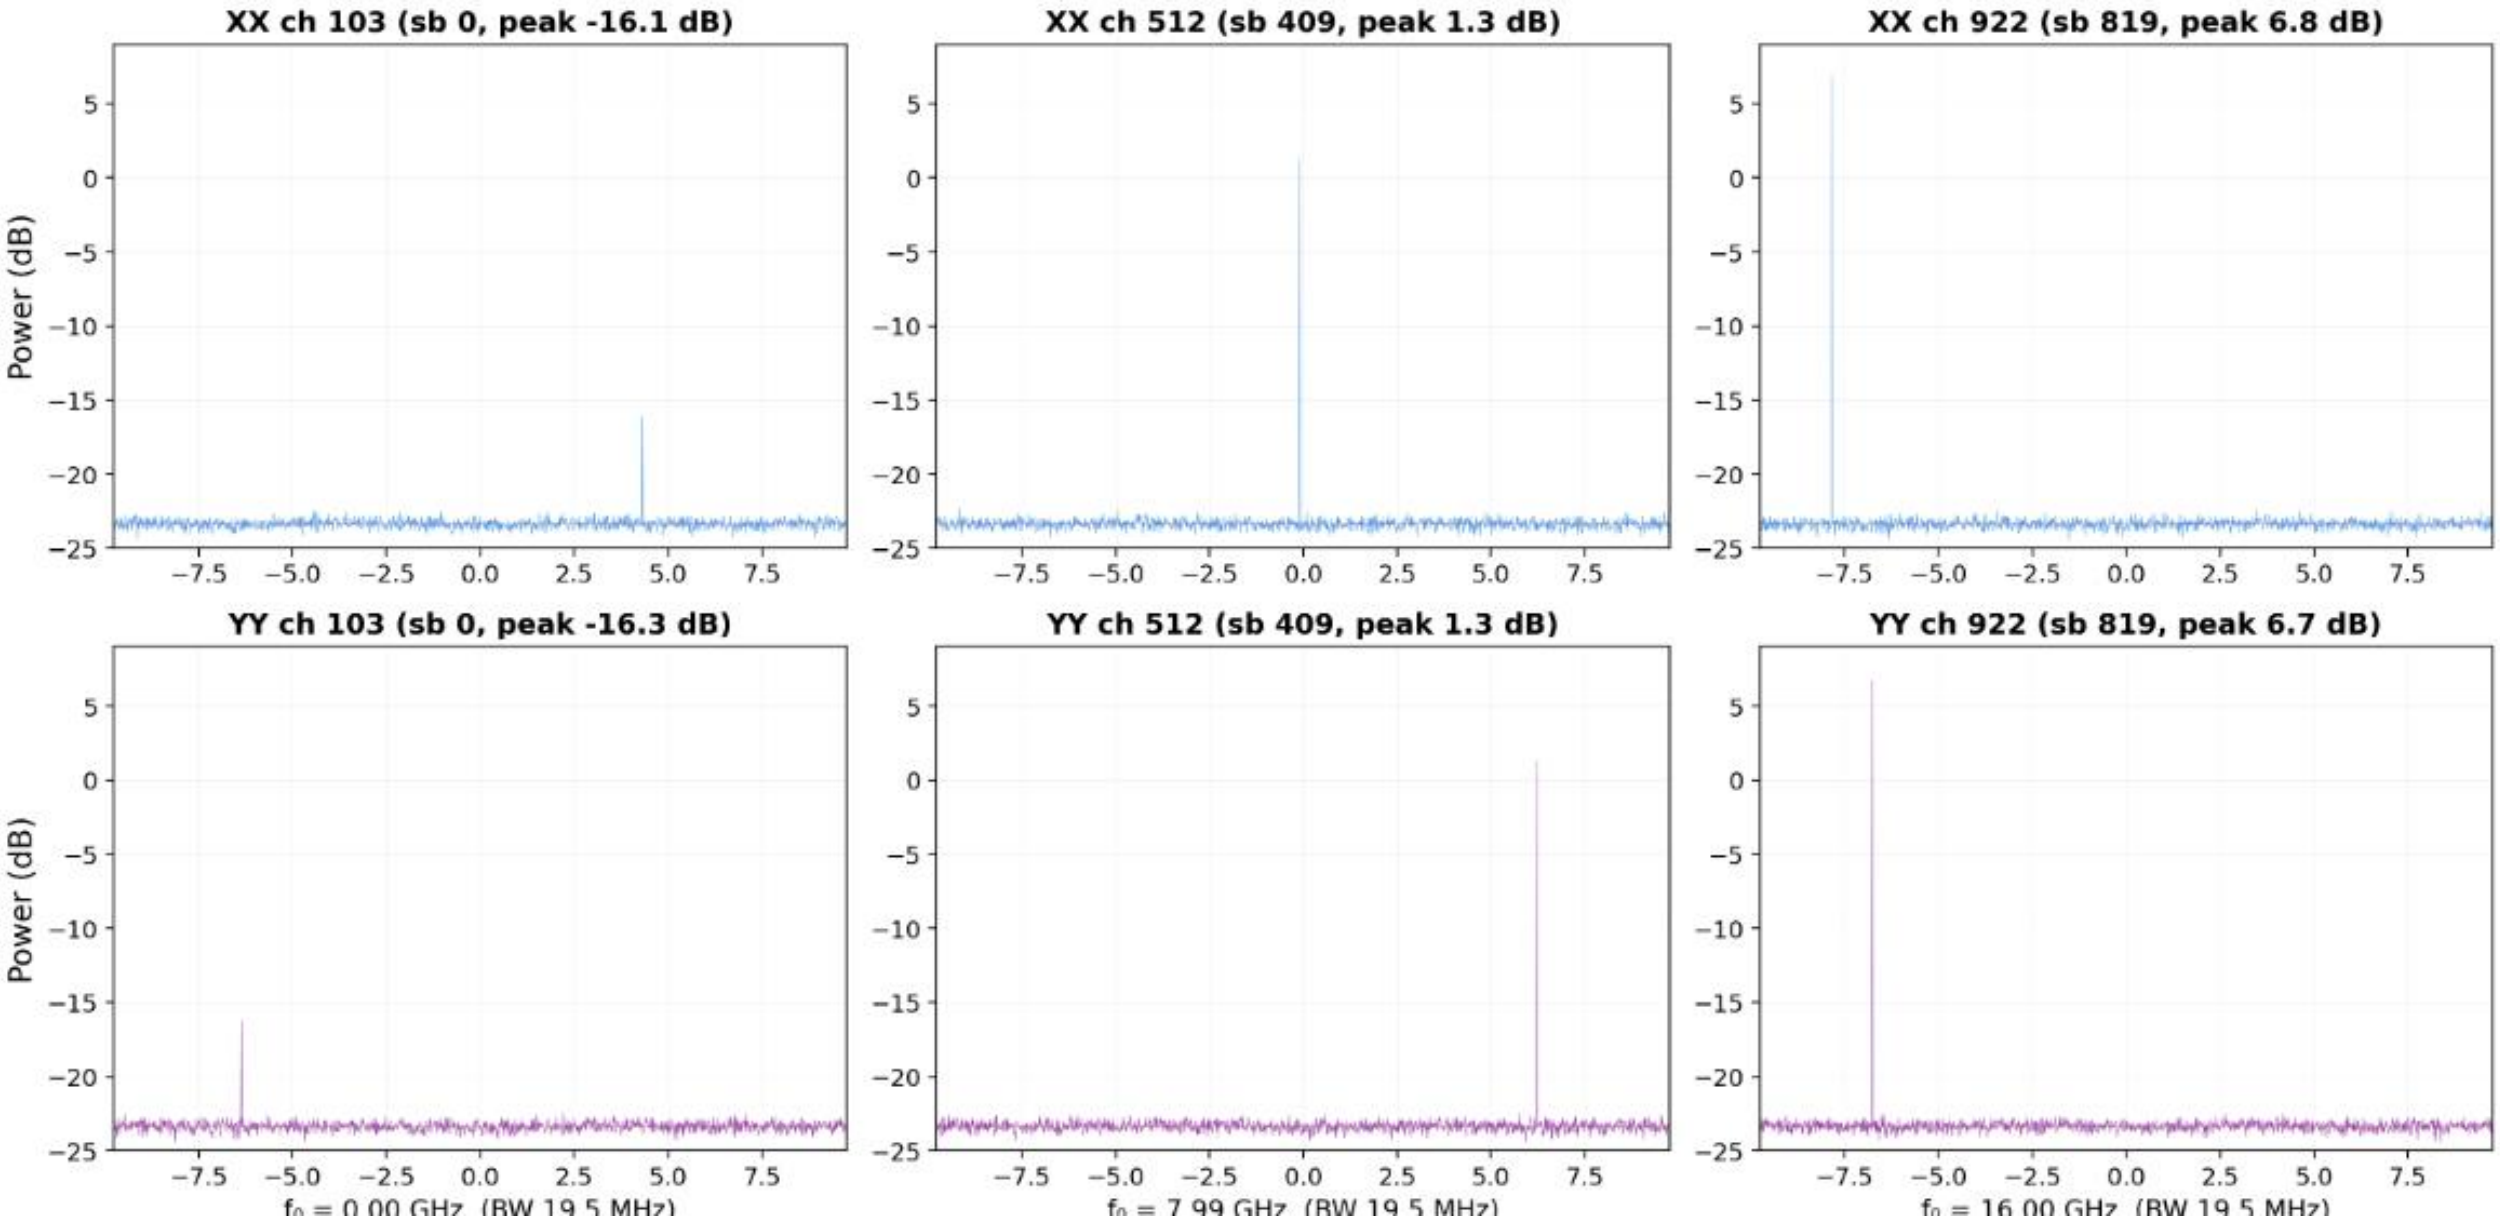

Summary

- TPGS shall process the data output rate from four ALMA TP antennas at **3.84 Tbps**, leveraging **400 GbE** and advanced GPUs to handle 32 GHz bandwidth per polarization.
- The software employs a three-stage design—**Coarse OSPFB**, **Tiled Transpose**, and **Fine CSPFB**—to transform 40 Gsps raw data into **13.5 kHz resolution spectra**.
- Performance verification confirms a clean noise floor and a final resolution of **13.49 kHz** across **1,187,360 fine channels**.
- Single-polarization processing currently meets the **16 ms real-time budget** on both GH200 (15.2 ms) and RTX 6000 (13.2 ms) hardware.
- Dual-polarization processing (22–23.8 ms) currently exceeds the time budget, necessitating next-generation **NVIDIA Rubin** GPUs and **PCIe Gen 6** for full real-time compliance.