Synthetic Aperture Radar (SAR) Back Projection Using AI Engines in AMD Versal™ Adaptive SoCs

Ken O'Neill – AMD Adaptive Embedded Computing Group
Mark Rollins – AMD Adaptive Embedded Computing Group
koneill@amd.com
mark.rollins@amd.com

ESA / IEEE European Data Handling and Processing Conference
October 2025

APPROVED FOR PUBLIC RELEASE

Agenda

- Synthetic Aperture Radar (SAR)
  - Background and introduction to SAR
  - Back projection processing

- Example of SAR back-projection processing in AMD Versal™ adaptive SoC platform
  - Design example methodology
  - System parameters
  - Implementation details
  - Results
  - Expansion in scope

- AMD XQR Versal Adaptive SoC for space applications
  - Qualification and reliability
  - Radiation effects

Synthetic Aperture Radar (SAR) – Background and Introduction

- Creates high resolution radar images from lower resolution radar samples limited by physical antenna size
  - High resolution radar imaging using conventional techniques requires a very large and expensive antenna
- Principle of SAR operation
  - The SAR platform passes over an object, illuminating it with many radar pulses from a smaller, cheaper antenna
  - Many returns are sampled, allowing the SAR platform to gather data multiple times from several viewing angles
- Traditional approaches to SAR require raw radar returns to be transmitted to ground for processing
  - Communication bottleneck limits the amount of imagery that can be gathered
  - Real-time processing on the SAR platform allows transmission of finished images, avoids communication bottleneck

In this graphic, point P is illuminated many times as the platform passes from point X_1 to X_2. Each time it is illuminated, it reflects radar energy ("backscatter") which is sampled by the SAR platform.

Image Credit: NASA SAR Handbook

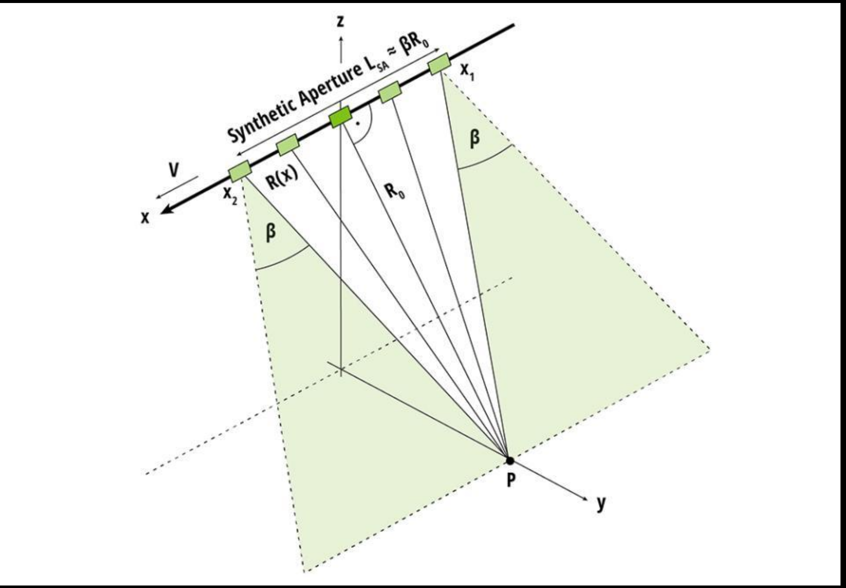

SAR Background – Range Dimension

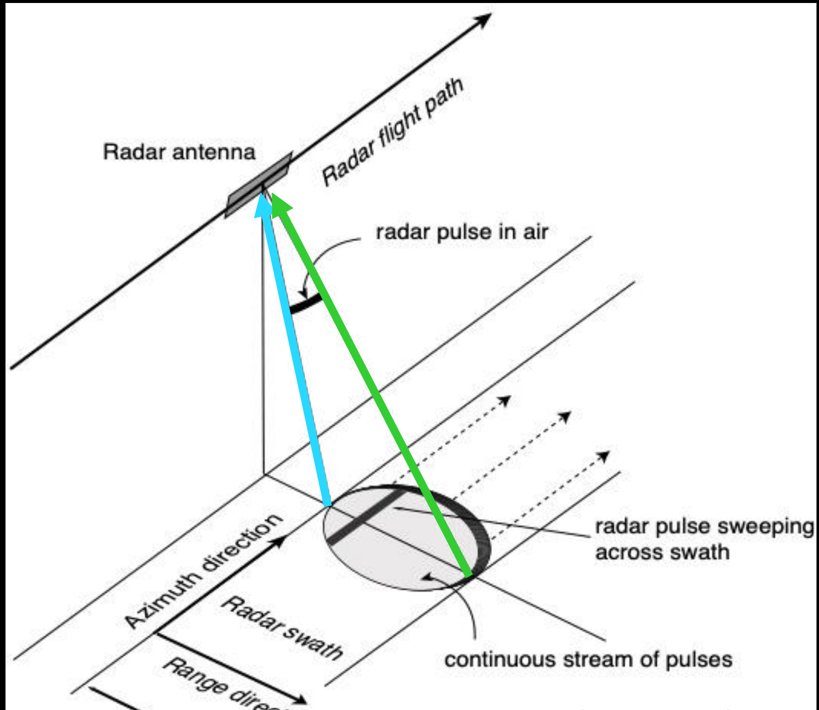

Image Credit: NASA

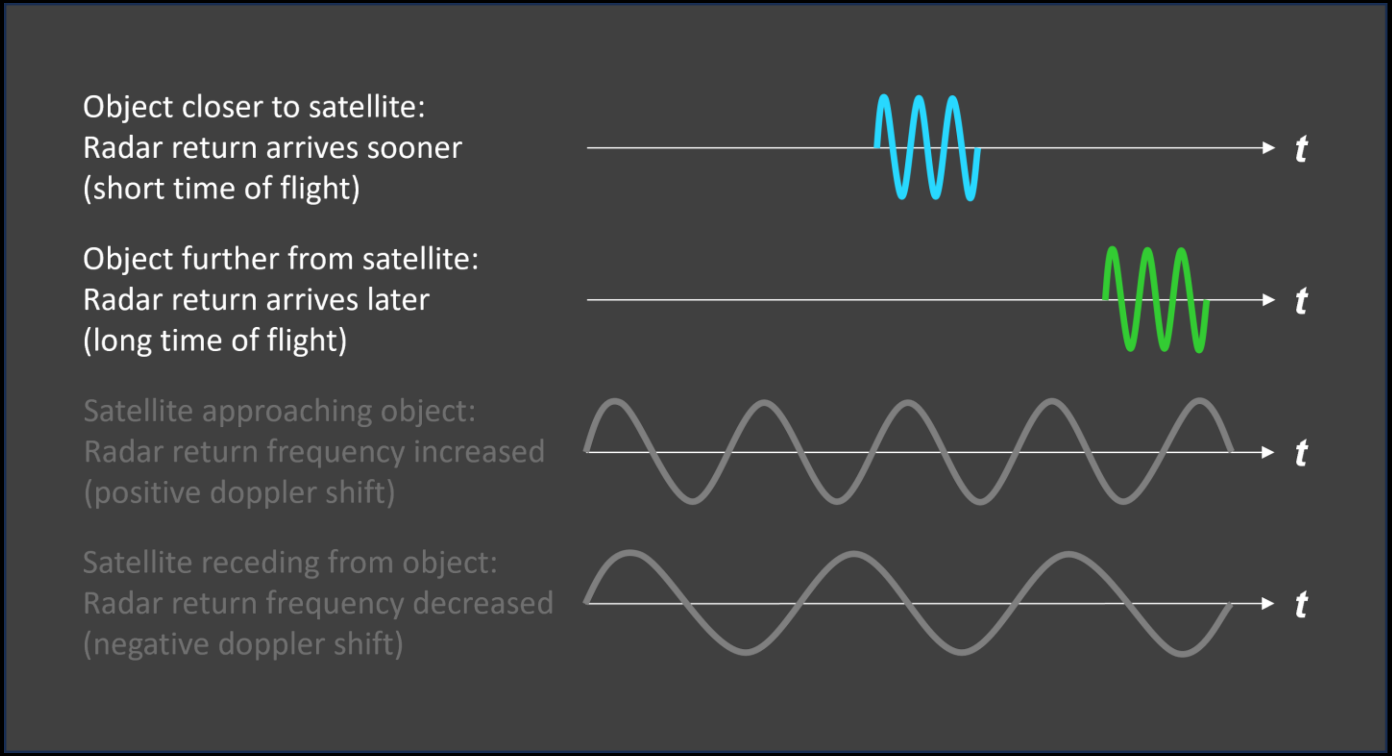

- Radar returns are sampled and stored for processing
- Reflections from near objects arrive sooner than reflections from distant objects

Acquiring Pulse Returns from a Single Pulse

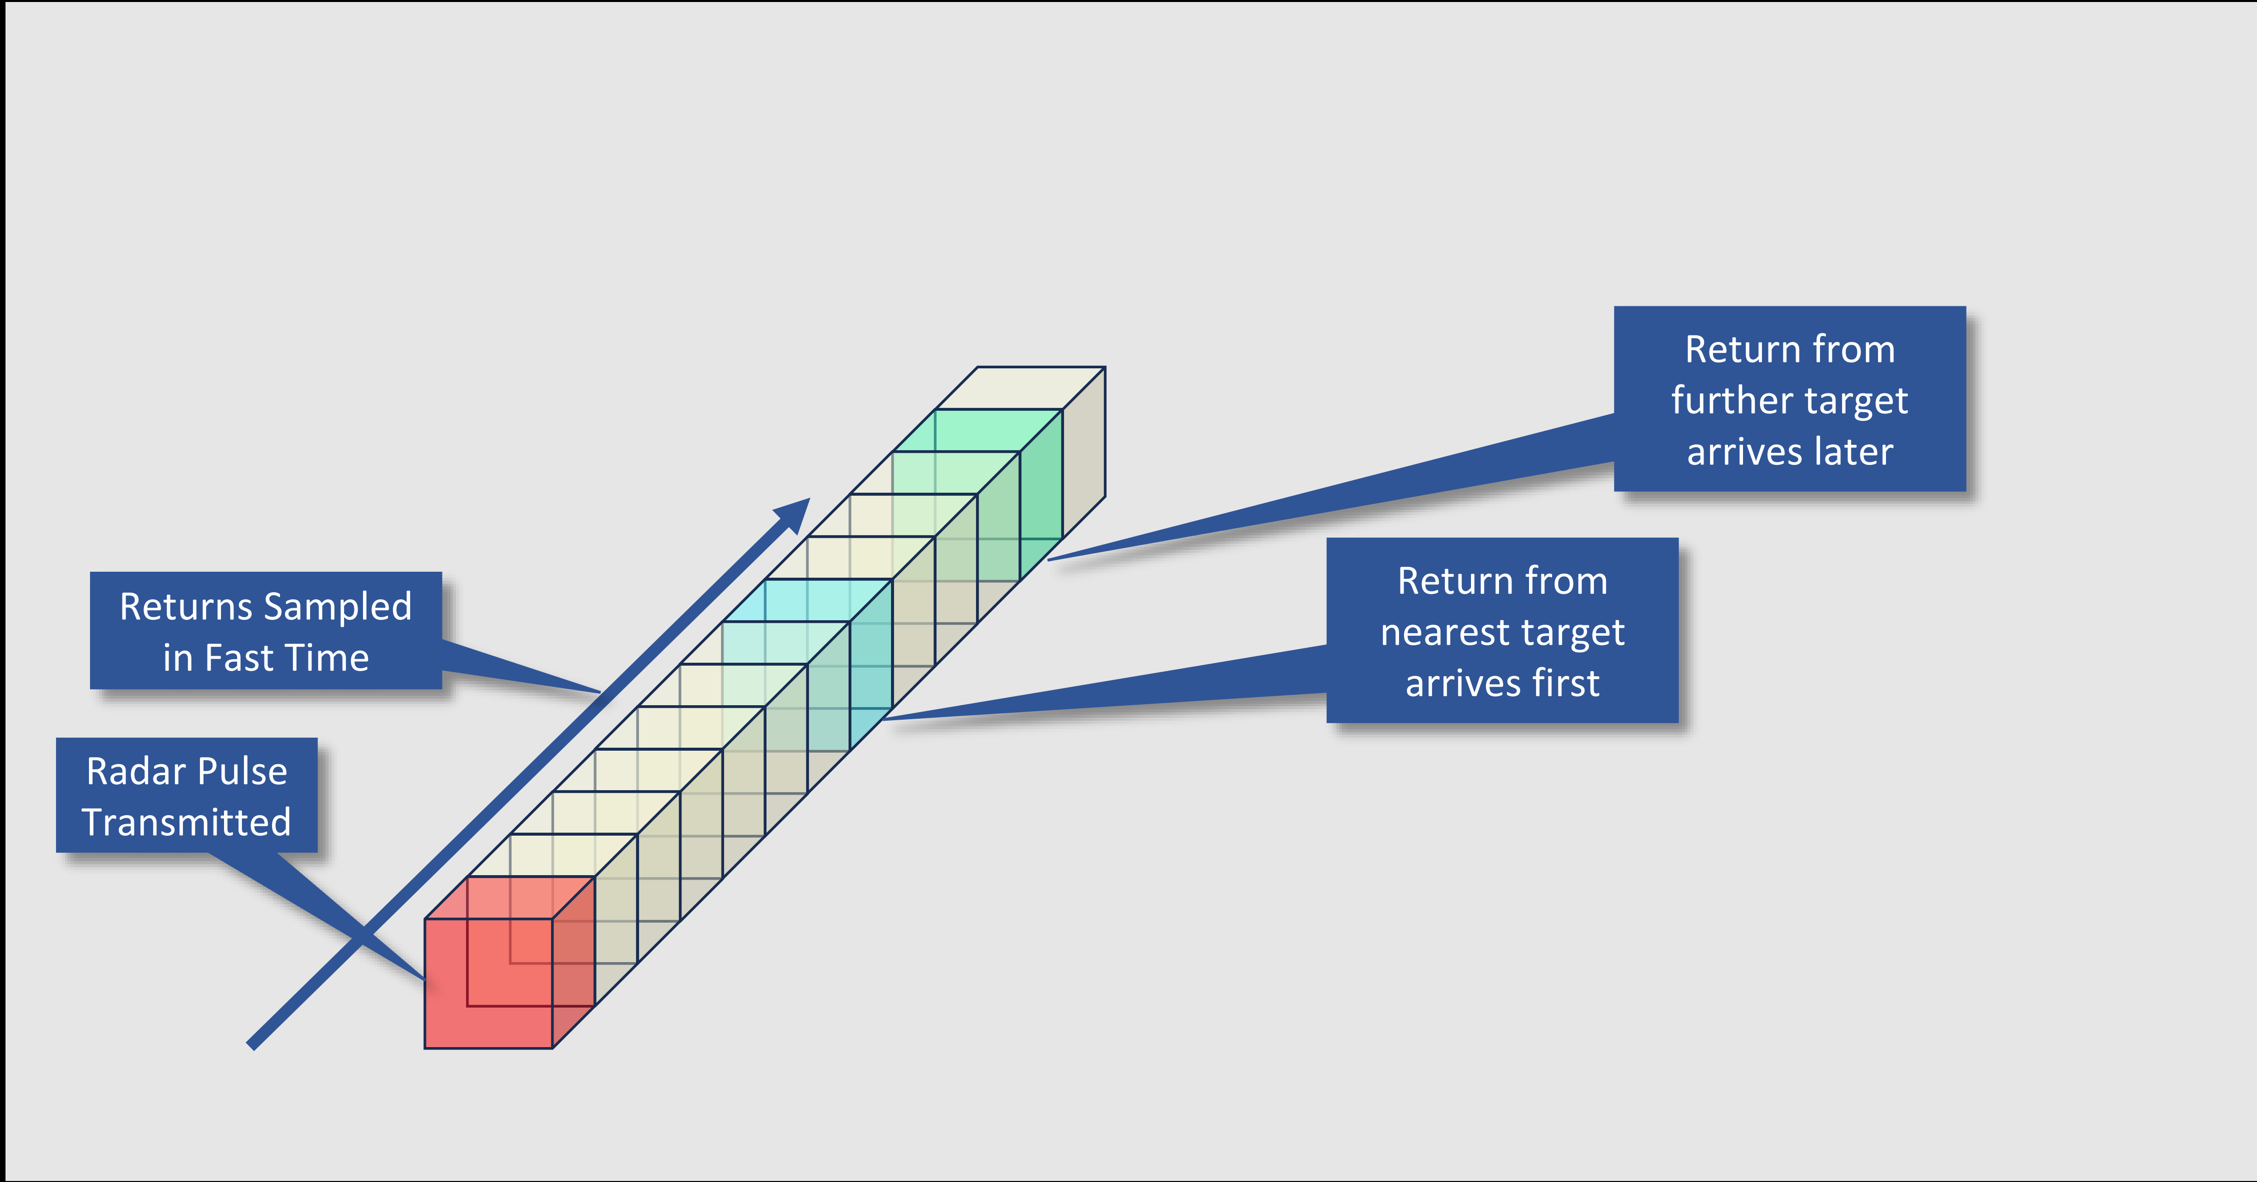

SAR Background – Doppler Dimension (Azimuth)

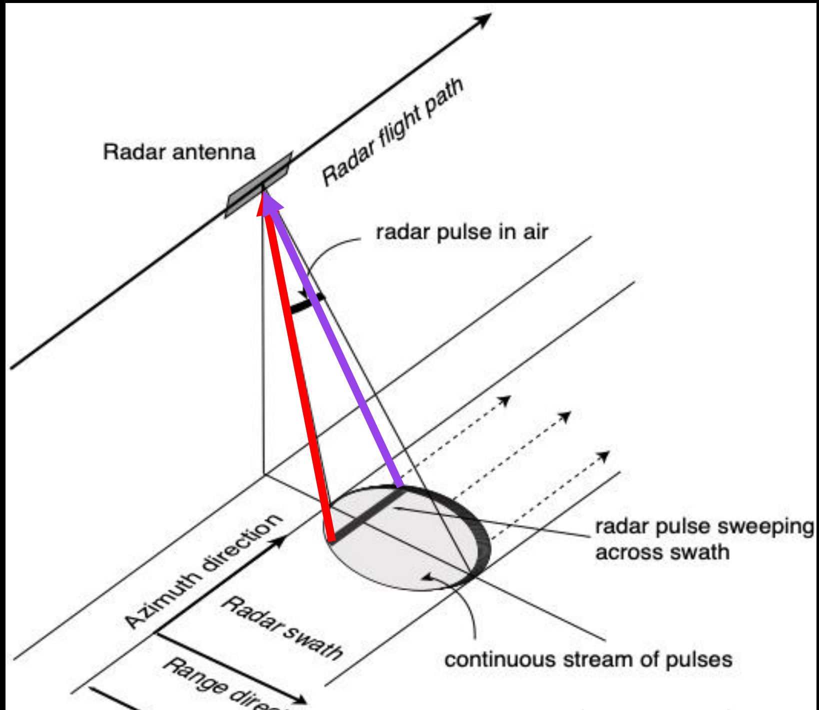

Image Credit: NASA

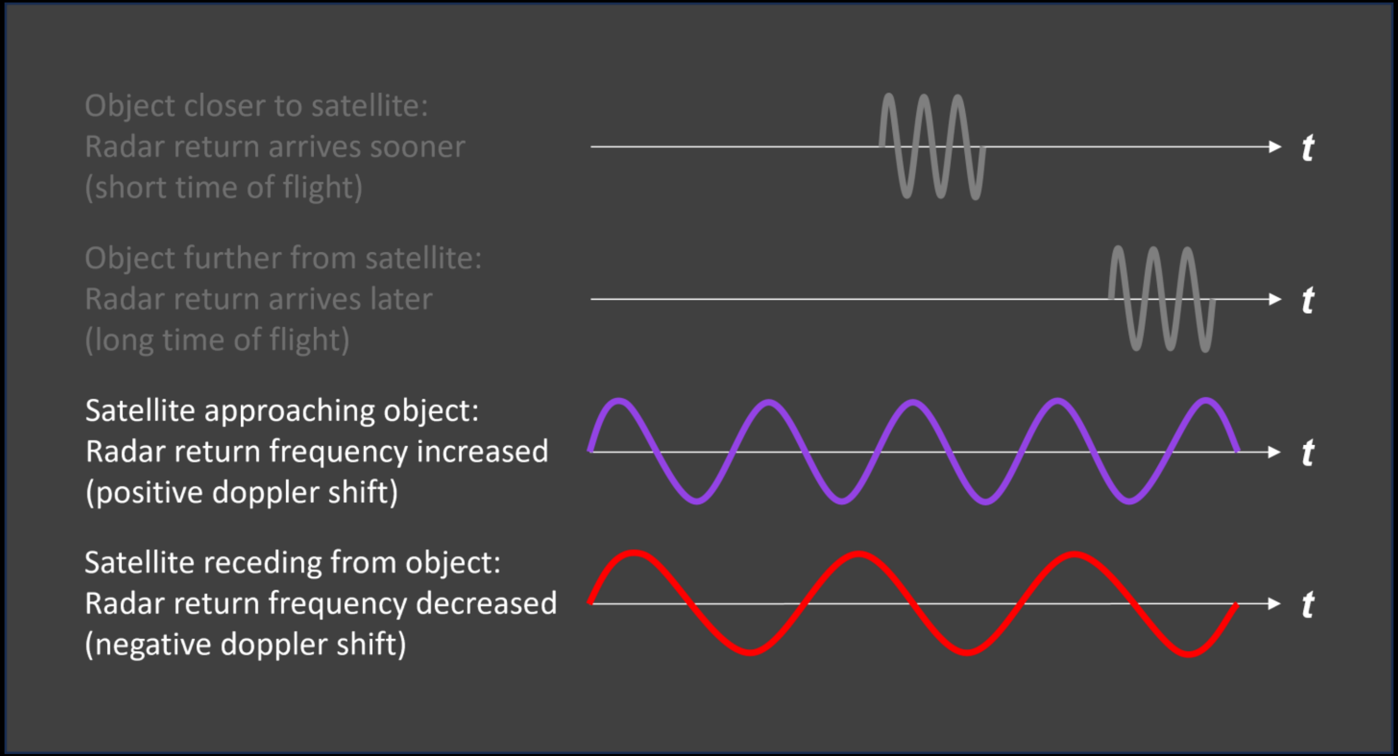

- Reflections from objects ahead of the satellite in the azimuth direction (direction of travel) will contain a positive doppler shift, which appears as an increase in carrier frequency
- Reflections from objects behind the satellite in the azimuth direction (direction of travel) will contain a negative doppler shift, which appears as a decrease in carrier frequency

AMD

Acquiring Pulse Returns from Multiple Pulses

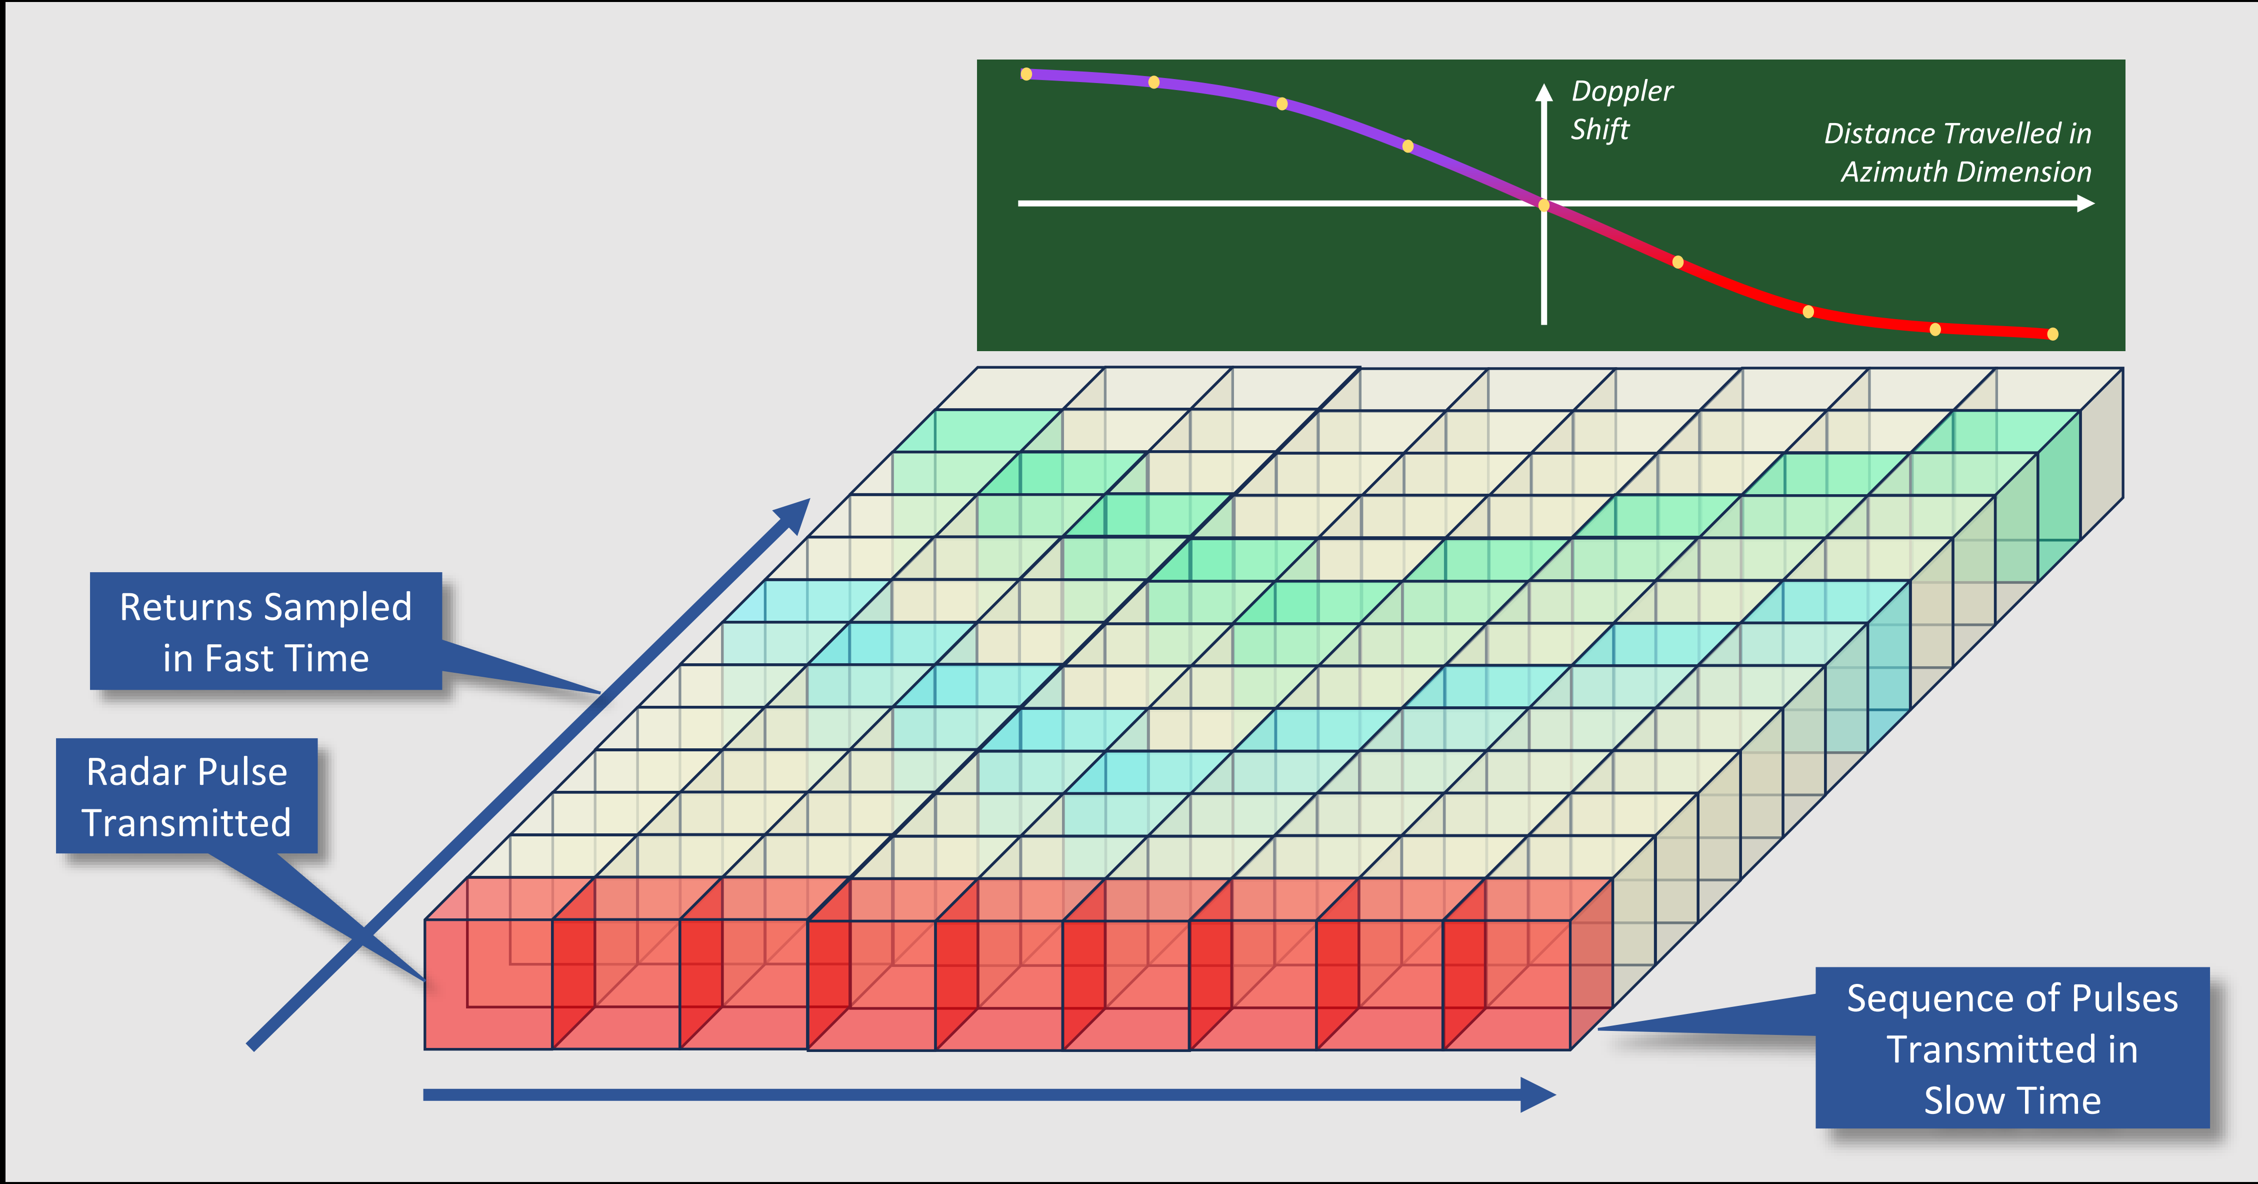

Principles, Advantages and Disadvantages of Back-Projection

- Principle of operation
  - Radar pulse returns (echos) are captured over time
  - Returns from multiple pulses are captured in a data array
  - For each pixel in the finished image, returns are adjusted for phase shift and added coherently
- Advantages and disadvantages of back-projection
  - Can provide higher quality images with greater levels of resolution that other SAR processing algorithms
  - Slower and more computationally intensive than other methods such as Range Doppler processing

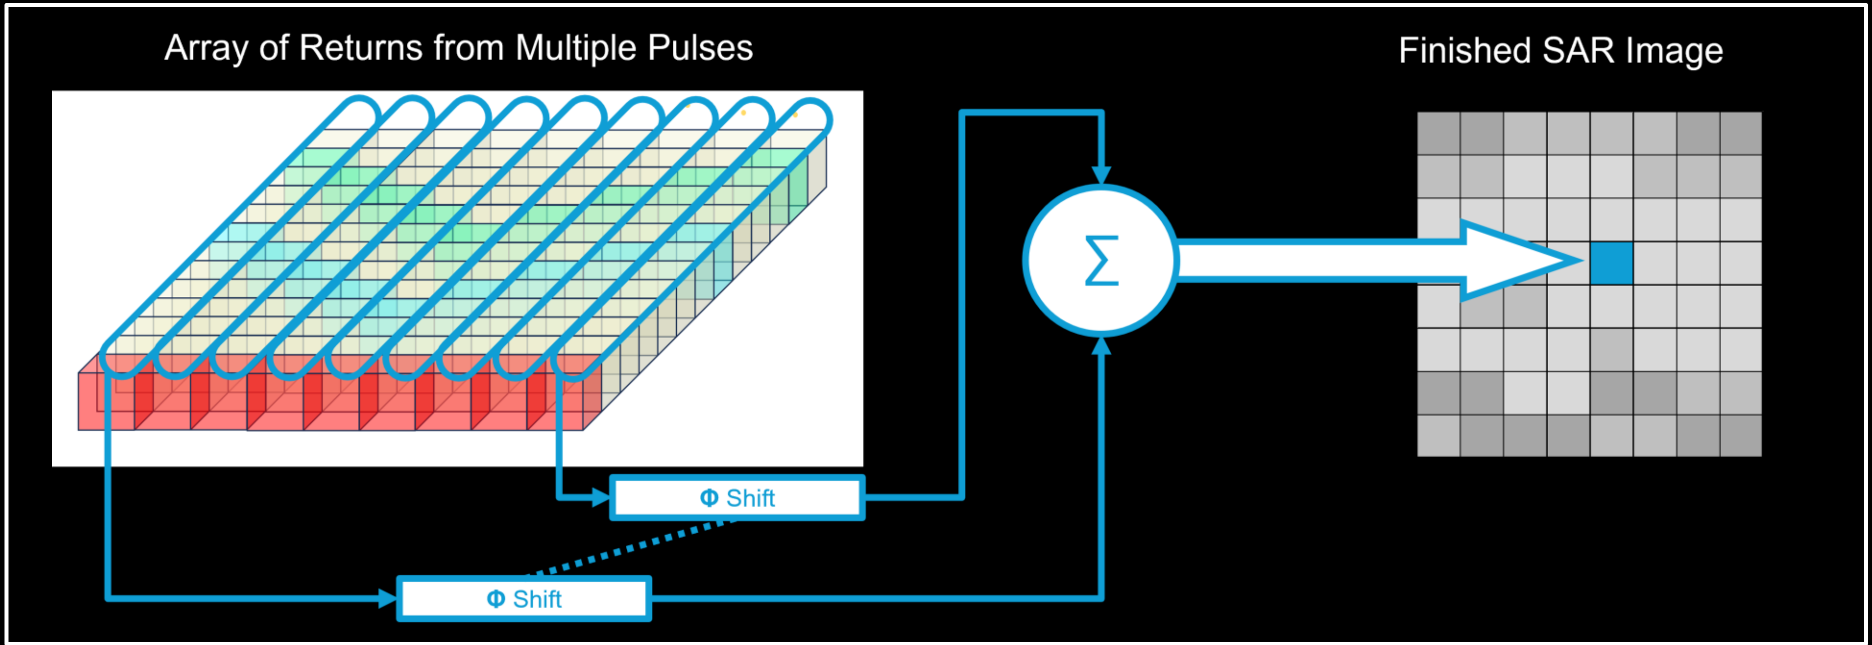

Back-Projection SAR Processing Design

- Design approach
  - Use SAR back-projection model for MATLAB® (Gorham and Moore, 2010); review simulated performance in Simulink
  - Implement in AMD Versal™ VC1902 adaptive SoC, assess utilization and performance, compare with simulation

Parameter	Value	Notes
Image Width × Height	512 × 512 Pixels	Square image
Number of Pulses	586 Pulses	For 5 azimuth angles
Target Throughput	1 GOP/sec	Rate of per-pixel back projection operations
IFFT Transform Size	2048 Points	Based on system model
Final Frame Rate	6.5 fps	All pulses accumulated
Per-Pulse Frame Rate	3820 fps	For a single pulse
Image Storage in DDR	2 MB	Assume 8B per pixel
IFFT Transform rate	3820 Hz	One transform per radar pulse
IFFT Sampling rate	8 Msps	Assume streaming solution
Total Number of AI Engines	TBD	Requires prototyping

Implementation of Back-Projection in AI Engines

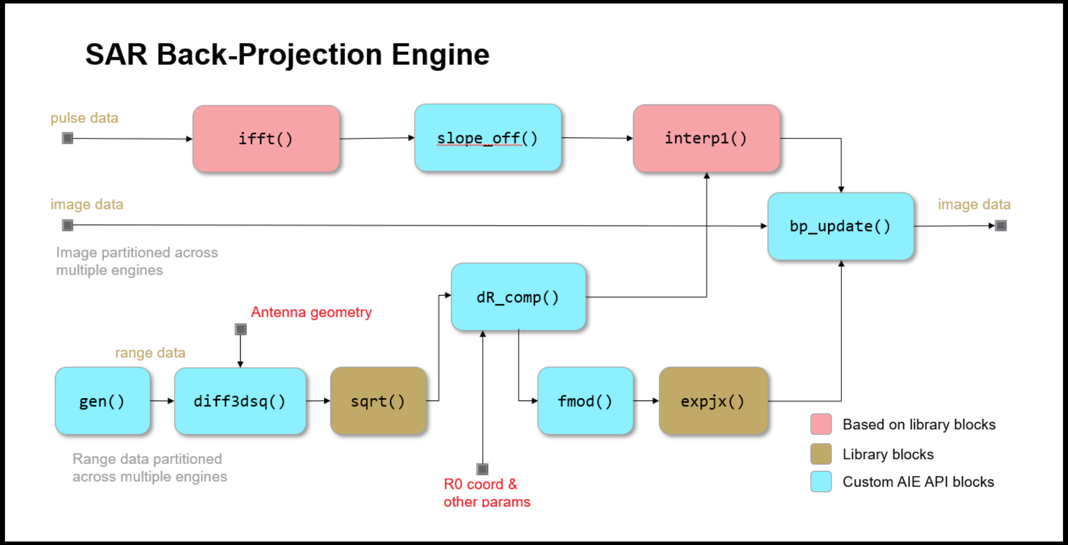

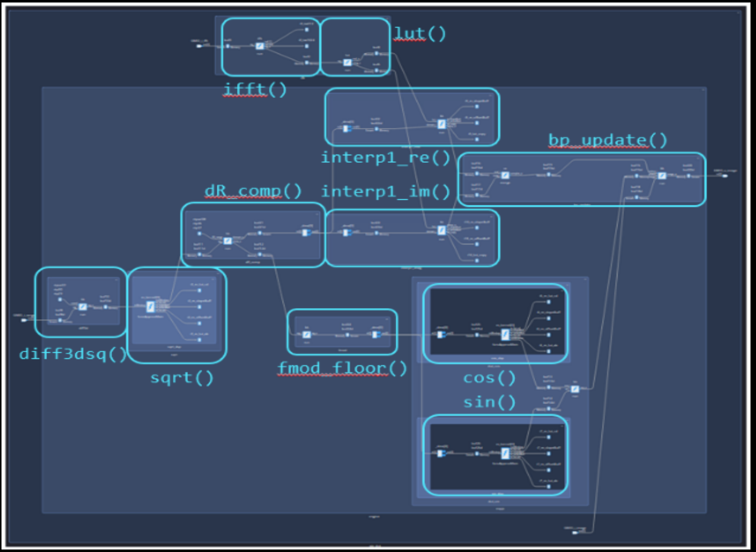

- Library blocks
  - sqrt() – compute square root of squared distance
  - expjx() – compute complex exponential function
- Modified library blocks
  - ifft() – inverse fast fourier transform
  - interpl() – interpolate differential distance
- Custom AIE blocks
  - diff3dsq() – compute 3D squared distance
  - slope_off() – compute slope and offset
  - dR_comp() – compute difference in range between target and scene center
  - fmod() – reduce input argument for exp()
  - bp_update() – update the SAR image using phase correction and distance terms
- Using single-precision floating point data type

AI Engine

Versal™ Adaptive SoC

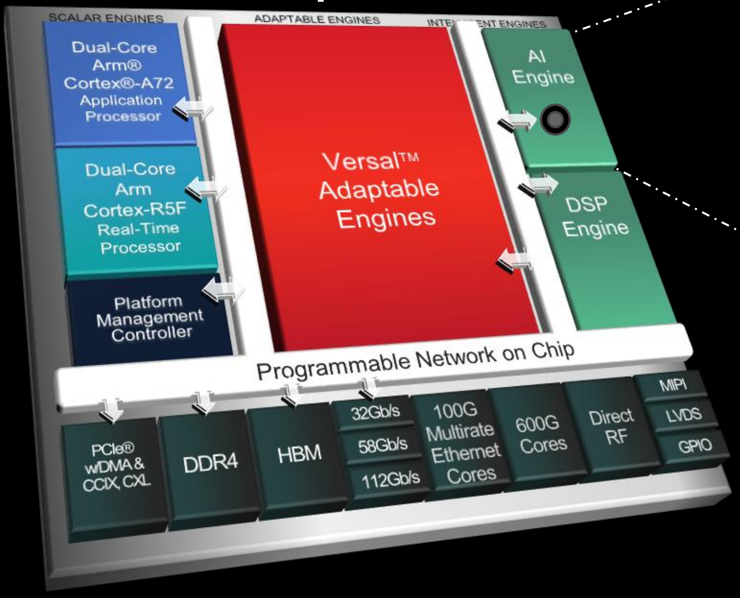

AI Engine Array

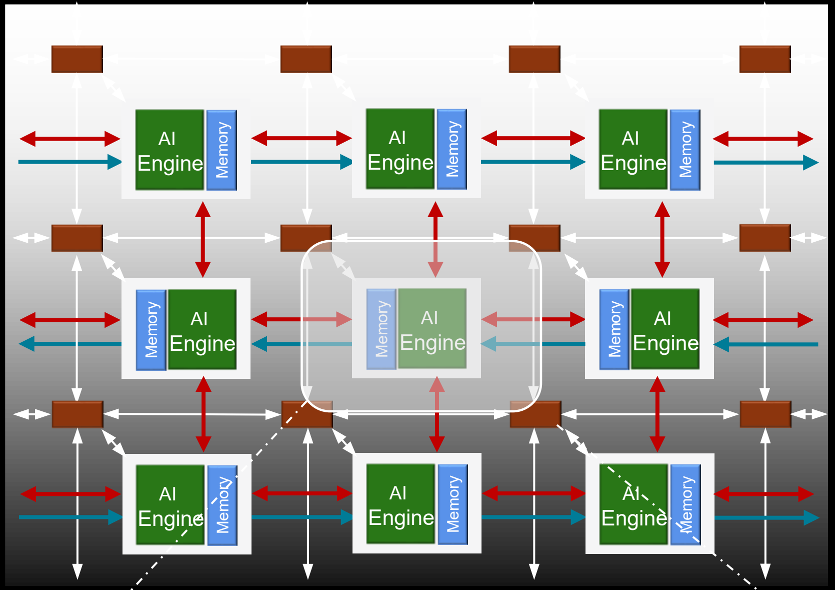

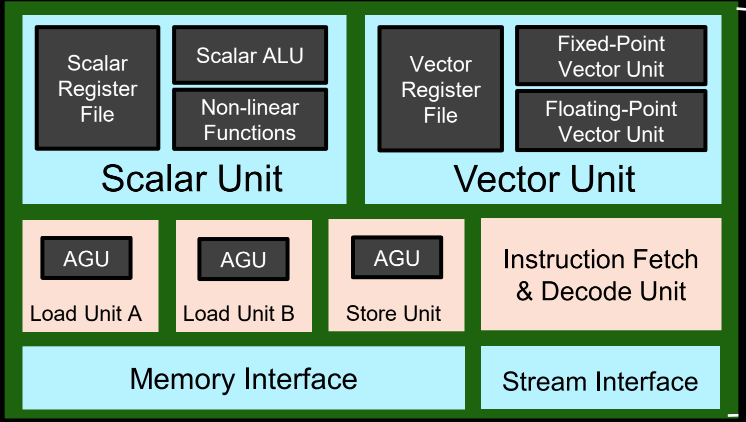

AI Engine

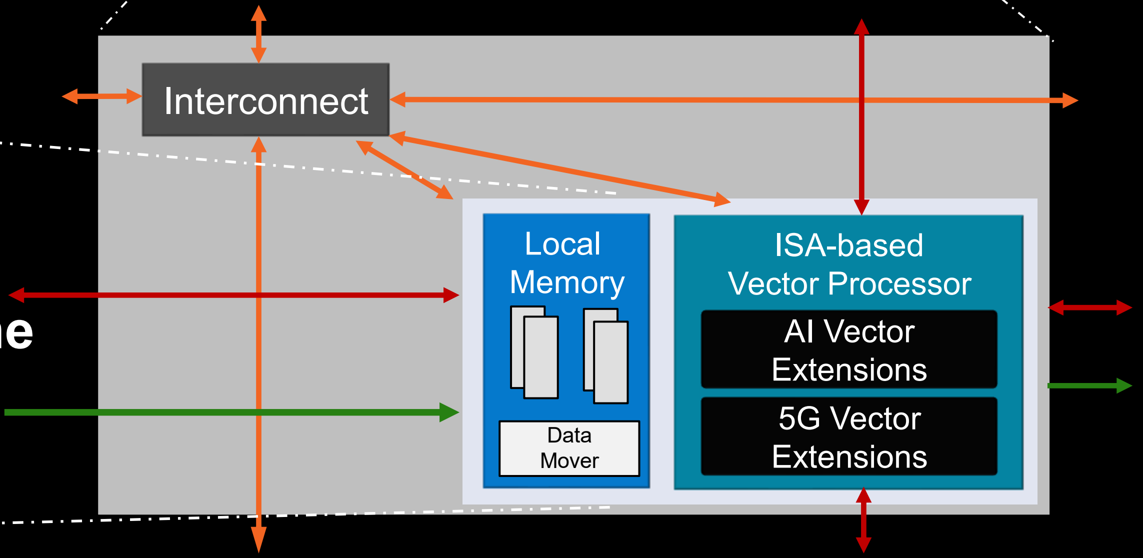

Memory Interface
Stream Interface
Cascade Interface

AMD

Device Resources and Results in AMD Versal™ Adaptive SoC

- Design targeted to AMD XCVC1902 Versal™ AI Core adaptive SOC
  - Ample resources for the design and can support expansion of the design
  - Space-qualified version is available for space-based SAR applications
- Implemented on AMD VCK190 evaluation board (uses XCVC1902)

Tiles used for Kernels, Buffers or Nets	30 of 400 (7.5%)
Tiles used for AI Engine Kernels	14 of 400 (3.5%)
Tiles used for Buffers	26 of 400 (6.5%)
Tiles used for Stream Interconnect	17 of 400 (4.3%)
GMIO Input Channel Usage	1 of 32
GMIO Output Channel usage	0 of 32
PLIO Input Channel Usage	1 of 312
PLIO Output Channel Usage	1 of 234
DMA FIFO Buffers	0
Interface Channels used for ADF Input / Output	3
Interface Channels used for Trace Data	0

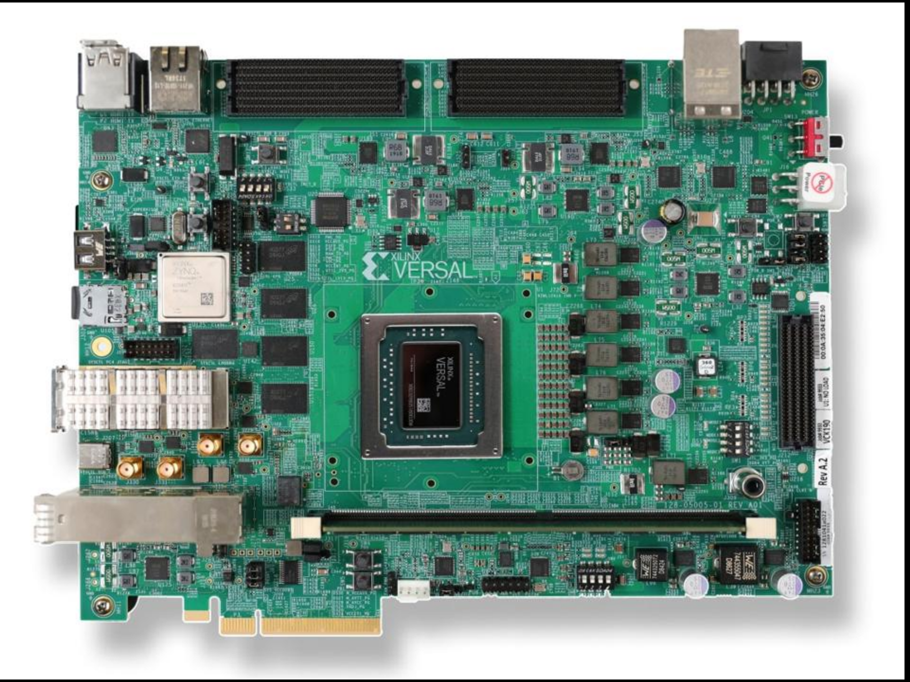

Performance in AMD Versal™ Adaptive SoC

- Frame rate 2.6 fps, 512 × 512 pixels per frame
- Latency 390 msec, 586 pulses
- Tested using USAF “GOTCHA” SAR data set

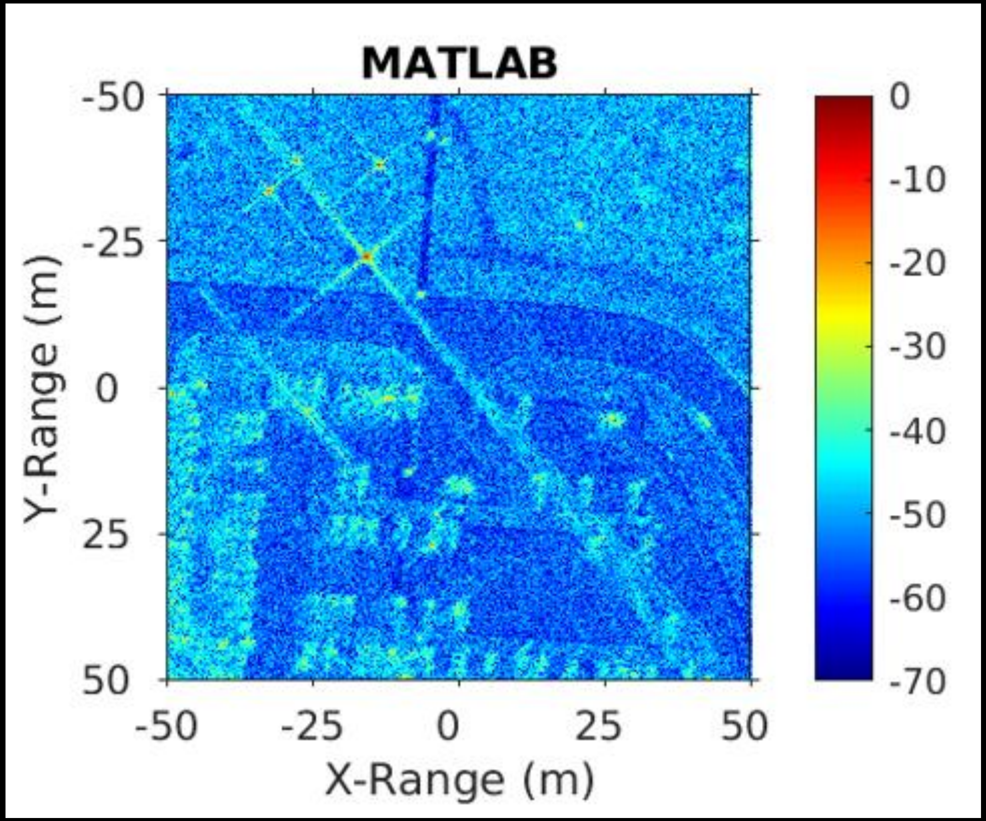

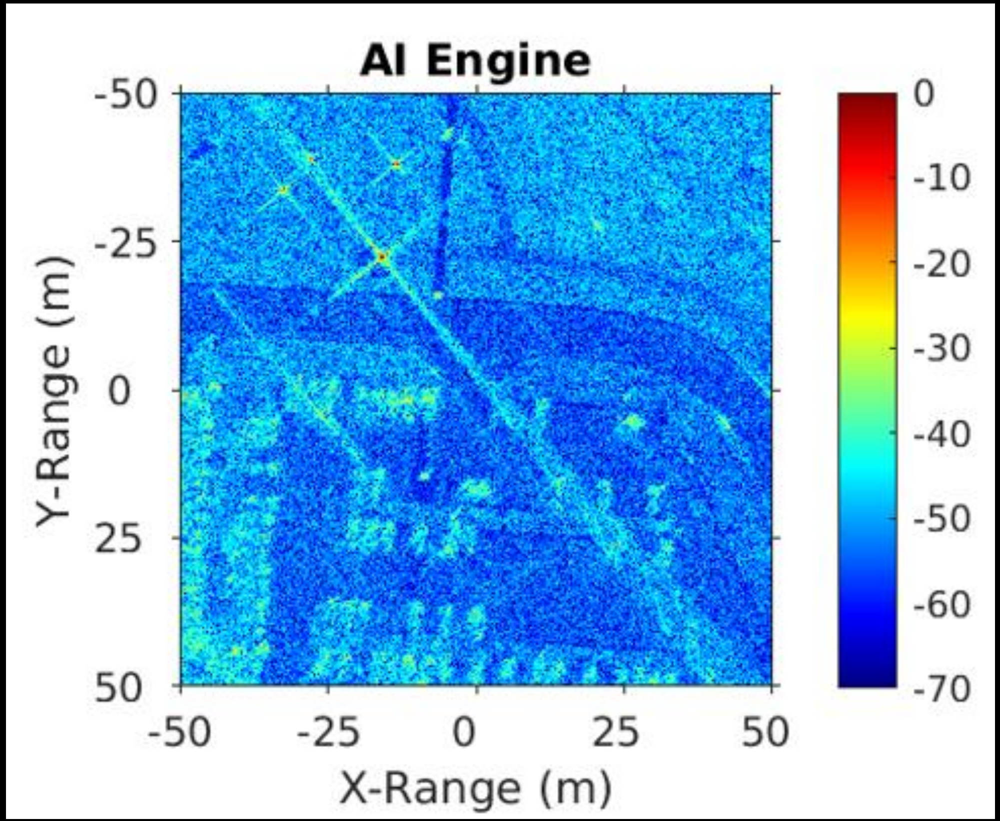

AMD

Quality of Results in AMD Versal™ Adaptive SoC

- Structural Similarity Index Measure (SSIM) = 0.9965
- Pixel error peak signal to noise (PSNR) = 74.1 dB

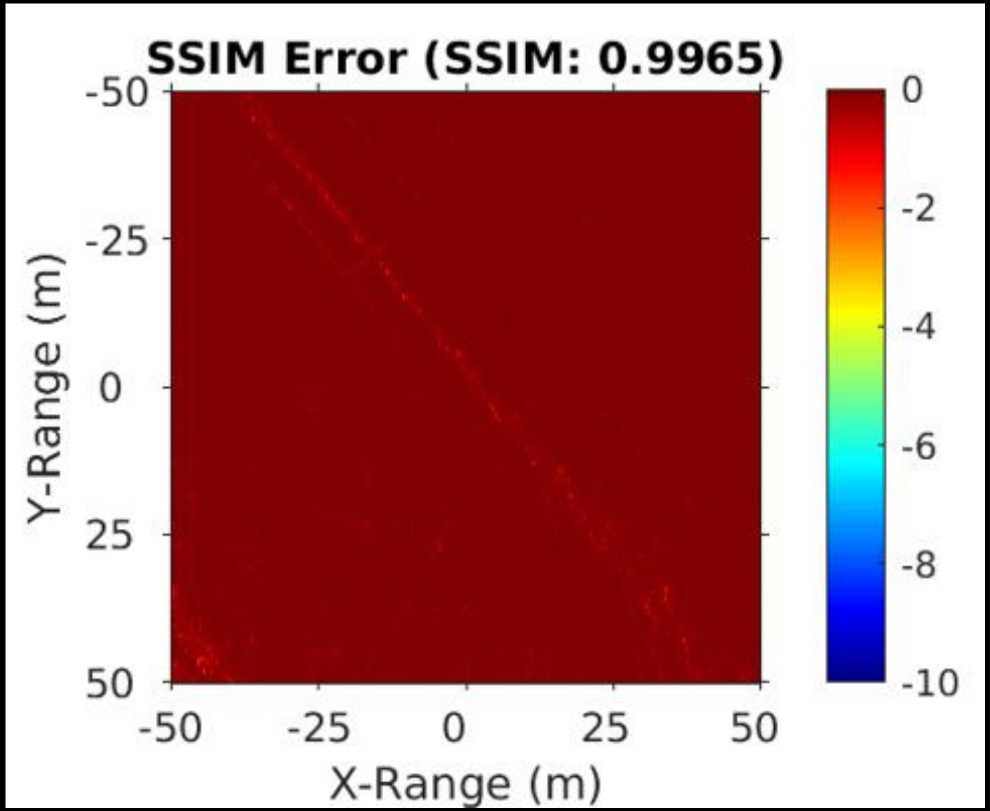

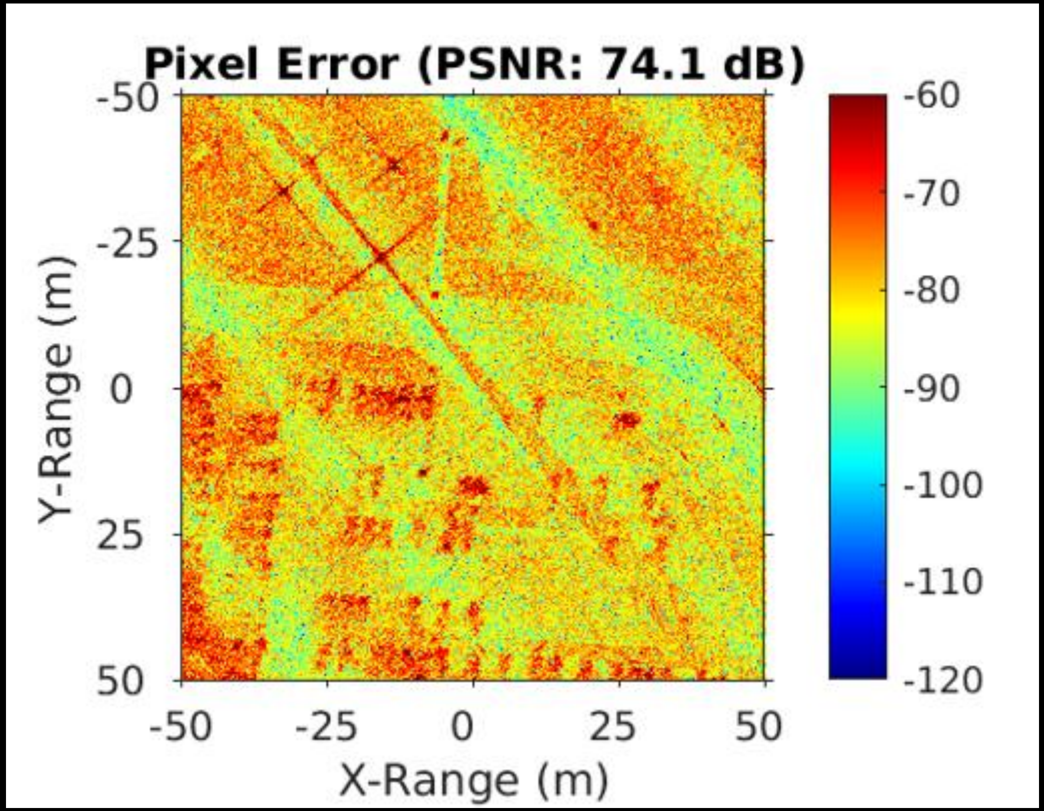

Expansion – Multiple Engine Design, Higher Throughput

- Eight copy step-and-repeat of original design to accomplish increase in throughput of approx. 8X
- Frame rate increased from 2.6 fps to 19.3 fps, increase of 7.42X
- Optimization is possible with code restructuring and modifications to memory architecture
-

Tiles used for Kernels, Buffers or Nets	193 of 400 (48.3%)
Tiles used for AI Engine Kernels	112 of 400 (28.0%)
Tiles used for Buffers	208 of 400 (52.0%)
Tiles used for Stream Interconnect	92 of 400 (23.0%)
GMIO Input Channel Usage	8 of 32
GMIO Output Channel usage	0 of 32
PLIO Input Channel Usage	8 of 312
PLIO Output Channel Usage	8 of 234
DMA FIFO Buffers	0
Interface Channels used for ADF Input / Output	24
Interface Channels used for Trace Data	0

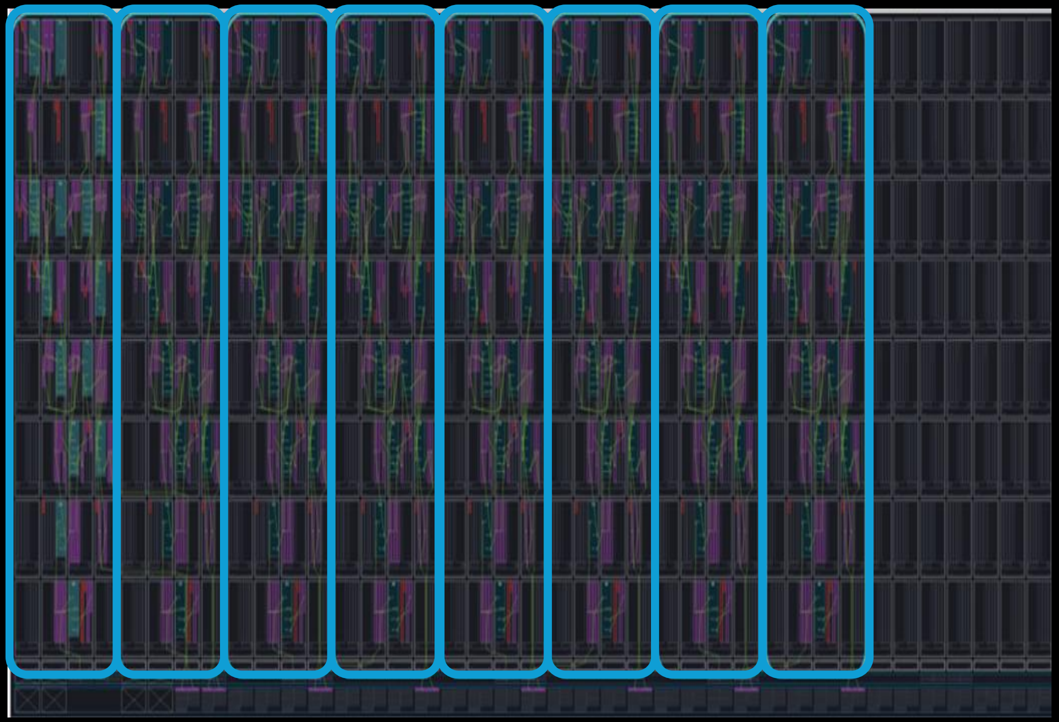

AMD

Conclusion

- Synthetic aperture radar (SAR) is a commonly-used technique to create high resolution radar images using an airborne or spaceborne platform
- Complex signal processing, filtering and FFTs are needed to perform SAR processing in real-time
- AI Engines in AMD Versal™ Adaptive SoC devices process SAR images in real-time on the radar platform, eliminating the need to transmit huge volumes of raw radar returns to the ground for processing
- Other work: Back-Projection processing is one of several methods for generating SAR images. AMD has also developed a Range / Doppler SAR processing reference design, which will be presented in the future.

- Ken O'Neill koneill@amd.com
- Mark Rollins mark.rollins@amd.com

AMD XQR Versal™ Adaptive SoC
For Space Applications

AMD Space Grade Products

Virtex™
XQR4QV

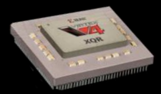

Shipping
Xilinx Class V
PDN Issued 10/21

Virtex™
XQR5QV

Shipping
Xilinx Class B, Class V

Kintex™
UltraScale
XQRKU060

Shipping
Xilinx Class B, Class Y

UltraScale+™
MPSoC/RFSoC

Shipping
Defense Grade*
*Potential for
proton based missions
(details shared with NDA)

XQR Versal™
Adaptive SoC
for Space

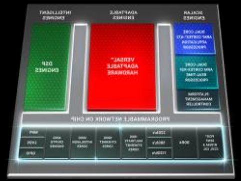

- Machine Learning
- Processor System
- High Density
- Internal Scrubber
- Ruggedized
- Defense Grade Chip Caps
- 7 Year Missions
- Xilinx Class B Qualified

XQR Versal
AI Core
XQRVC1902

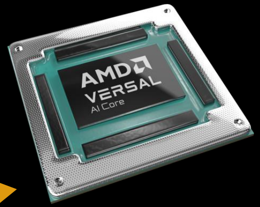

Class B Qualified,
Shipping Now

XQR Versal
AI Edge
XQRVE2302

Class B Qualified,
Shipping 4Q2025

First 7nm Programmable Device
for Space Applications

Exceptional Compute
Performance for
On-board Data Processing!

90nm | 65nm | 20nm | 16nm | 7nm

XQR Versal™ Adaptive SoC for Space Product Table

		XQRVC1902-1MSBVSRA2197 (AI Core)	XQRVE2302-1MSBSSRA784 (AI Edge)
Intelligent Engines	**AI Engine Tiles**	**400 (AIE)**	**34 (AIE-ML)**
	AI Engine Data Memory (Mb)	100	17
	AI-ML Shared Memory (Mb)	-	68
	DSP Engines	**1,968**	**464**
Adaptable Engines	System Logic Cells (K)	1,968	329
	6-Input LUTs	**899,840**	**150,272**
	NoC Master/NoC Slave Ports	28	5
	Distributed RAM (Mb)	27	4.6
Memory	Total Block RAM (Mb)	34	5.4
	UltraRAM (Mb)	130	43.6
	Accelerator RAM (Mb)	-	32
	Total PL Memory (Mb)	**191**	**86**
	DDR Memory Controllers	4	1
	DDR Bus Width	256	64
Scalar Engines	**Application Processing Unit**	**Dual-core Arm® Cortex®-A72, 48KB/32KB L1 Cache w/ECC 1 MB L2 Cache w/ECC**	
	Real-time Processing Unit	**Dual-core Arm Cortex-R5F, 32KB/32KB L1 Cache, and 256KB TCM w/ECC**	
	Memory	256KB On-Chip Memory w/ECC	
	Connectivity	Ethernet (x2); UART (x2); CAN-FD (x2) USB 2.0 (x1); SPI (x2); I2C (x2)	
Serial Transceivers	**GTx Transceivers**	**44 GTY (26.5625 Gb/s)**	**8 GTYP (26.5625 Gb/s)**
Integrated Protocol IP	CCIX & PCIe® w/DMA (CPM)	1 x Gen4x8, CCIX	-
	PCI Express	4 x Gen4x8	1 x Gen4x8
	Multirate Ethernet MAC	4	1
	Platform Management Controller	Boot, Security, Safety, Monitoring, High-Speed Debug, SEU Mitigation (XilSEM)	
Package	**Ruggedized Organic BGA**	**VSRA2197, 45mm x 45mm, 0.92mm pitch**	**SSRA784, 23mm x 23mm, 0.8mm pitch**
I/O		648 XPIO, 44 HDIO, 78 MIO, 44 GTY	216 XPIO, 22 HDIO, 78 MIO, 8 GTYP
Radiation Single Event Effects (SEE)	Proton and Heavy-Ion SEE Testing	NO SEL, 100% Correctable SEUs, Ultra-low SEFI	
Qualification and Availability		**B-flow Qualified, Shipping Now**	**B-flow Qualified, Shipping 4Q2025**

XQRVC1902 Datasheet DS946*, *XQRVE2302 Datasheet DS955

AMD

XQRVE2302 Versal™ AI Edge Adaptive SoC

- Versal AI Edge adaptive SoC XQRVE2302 is now qualified to B flow
  - Significantly lower power consumption than XQRVC1902
  - Significantly less board space than XQRVC1902
  - Second generation AI engines ("AIE-ML") have increased throughput, optimized for AI inferencing
  - Qualification completed, product shipping 4Q2025

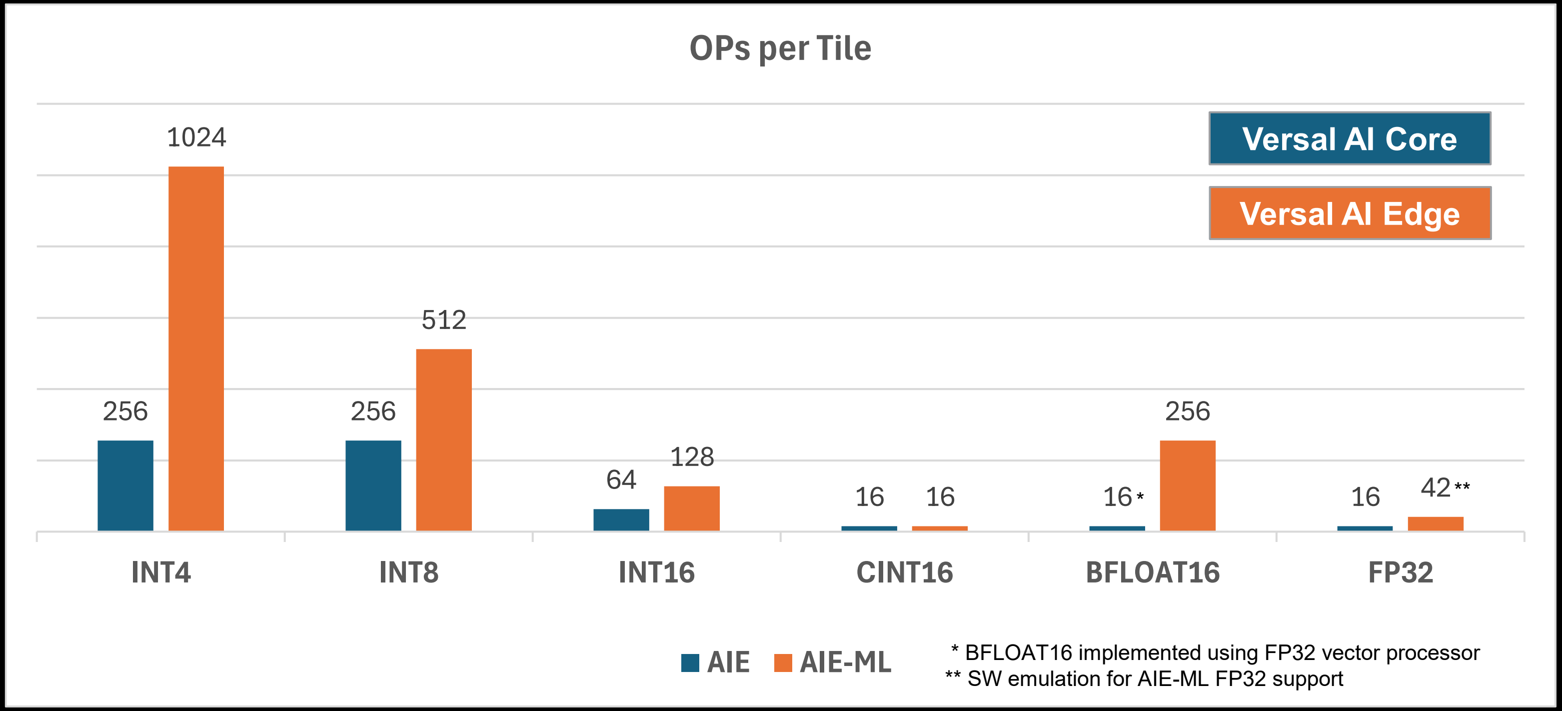

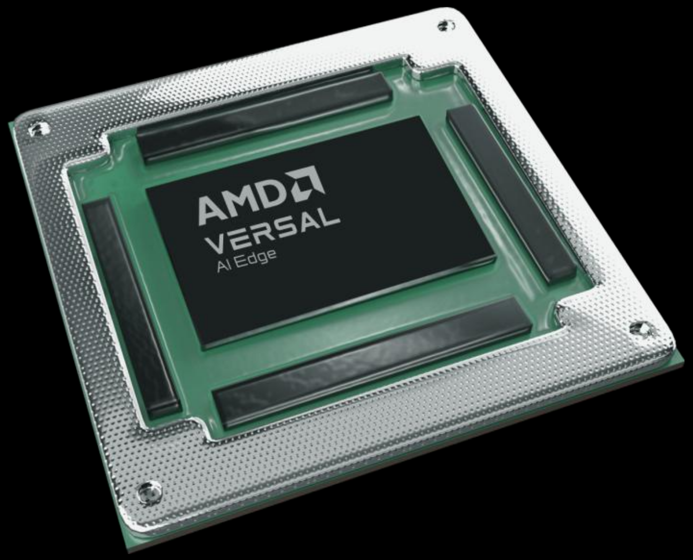

AMD

XQR Versal™ Adaptive SoC
Packaging, Qualification and Screening

Versal™ Adaptive SoC Class B Qualification Summary

- AMD has successfully completed our Class B qualification for the Versal XQRVC1902 device

Stress Test	MIL-STD-883 JEDEC reference	Conditions	Duration / Sample Size	Results
Prod. Burn-in	TM 1015	Dynamic, Tj = 125°C Vccmax	160 hrs.	Passed
Group A	TM 5005	Functional, AC and DC Parameters Test at -55°C, 25°C and 125°C	Test at -55°C, 25°C and 125°C	Passed
Group B	Various JEDEC	Assembly Monitors	✓	Passed
Group C [2]	TM 1005	Tj = 125°C, Vccmax	2 lots, 90 units total - 1000 hours 1 lot, 45 units – 10,000 hours	Passed
HTS [1]	TM 1008	Ta = 150°C	1000 hours 3 lots, 75 units total	Passed
THB [1]	JESD22-A101	85°C / 85% RH, Vccmax	1000 hours 3 lots, 75 units total	Passed
Temp Cycle [1]	TM 1010	B: -55°C / 125°C	1000 cycles 3 lots, 75 units total	Passed
Group D [1]	TM 5005	Sub-Groups 1,3,4,5	3 lots, 15 units / subgroup	Passed

(1) Units submitted to MSL-4 preconditioning prior to stressing

AMD Qualification for Versal™ XQRVE2302-SSRA784 (Class B)

- AMD has successfully completed our Class B qualification for the Versal XQRVE2302 device

Stress Test	MIL-STD-883 JEDEC reference	Conditions	Test Vehicle / Sample Size	Results
Prod. Burn-in	TM 1015	Dynamic, Tj = 125°C Vccmax. 160 hrs.	All units	Passed
Group A	TM 5005	Functional, AC and DC Parameters Test at -55°C, 25°C and 125°C	XQRVC1902-VSRA2197: 3 lots XQRVE2302-SSRA784: 1 lot	Passed
Group B	Various JEDEC	Assembly Monitors	XQRVC1902-VSRA2197: 3 lots XQRVE2302-SSRA784: 1 lot	Passed
Group C	TM 1005	Tj = 125°C, Vccmax, 1000 hrs.	XQRVC1902-VSRA2197: 3 lots, 135 units XQRVE2302-SSRA784: 1 lot, 45 units	Passed
HTS[1]	TM 1008	Ta = 150°C, 1000 hrs.	XQRVC1902-VSRA2197: 3 lots, 75 units	Passed
THB[1]	JESD22-A101	85°C / 85% RH, Vccmax, 1000 hrs.	XQRVC1902-VSRA2197: 3 lots, 75 units	Passed
Temp Cycle[1]	TM 1010	B: -55°C / 125°C, 1000 cycles	XQRVC1902-VSRA2197: 3 lots, 75 units XQRVE2302-SSRA784: 1 lot, 30 units	Passed
Group D[1]	TM 5005	Sub-Groups 1,3,4,5	XQRVC1902-VSRA2197: 3 lots, 15 units XQRVE2302-SSRA784: 1 lot, 15 units	Passed

(1) Units submitted to MSL-4 preconditioning prior to stressing

In support of the technology qualification, the XQRVC1902 device has passed 10,000 hours of Group C stressing

XQR Versal™ Adaptive SoC Radiation Effects

Versal™ Adaptive SoC Radiation Effects Summary

	Protons (2 – 105 MeV) Low Earth Orbit, 500 km, 20° inclination			Heavy-ions (1 - 80 MeV·cm²/mg) Geosynchronous Earth Orbit			TID (gamma)
	CRAM SEU (upset/bit/day)	SEL	SEFI (events/device/year)	CRAM SEU (upset/bit/day)	SEL	SEFI (events/device/year)	
Observed Rates	**3.5×10^{-9}**	**ZERO** events observed	PS: **1.3** XilSEM: **ZERO** AIE: **1.5** GT: 2025	**6.5×10^{-12}**	**ZERO** events observed	PS: **0.16** and XilSEM: **4.9×10^{-3}** AIE: 2025 GT Quad: **1×10^{-3}**	**PASS 120 KRad(Si)**
Comments	Proton energy: 64-400MeV Environment: 1×10^{12} p/cm² at 125°C			Ion Energy: 1-80 MeV·cm²/mg Environment: 1×10^{7} per ion/cm² at 125°C			<18 Krad/min

Estimates based on CREME96 AP8-Max; 500km and GEO models

- DUTs: Versal 7nm VC1902, 20 parts from 5 wafer lots to account for lot-to-lot variation
- ZERO SEL events in maximum V_{CC} and junction temperature conditions at LET up to 80 MeV·cm²/mg
- ZERO uncorrectable Configuration RAM (CRAM) events in LEO and GEO
  - Configuration RAM protected by EDAC and interleaving
- Robust XilSEM internal scrubber SEFI rate may eliminate need for on-board scrubber in space flight
  - Reference AMD user guides UG643 and PG352 for XilSEM scrubbing operation and cycle time
- AMD has published Versal SEE results at SEE/MAPLD 2022, NSREC 2022, 2023 and 2024, RADECS 2022 and 2023
  - Check AMD Space Lounge for new reports, links to conference papers and updated content

AI/ML Radiation Induced Datapath Error Signatures

Example: Misclassification
(Accuracy Degradation)

Actual Image[(1)]

Predicted Image

Image: ILSVRC2012_val_00000383
Model: resnet-18 (SAT)

Golden Model Prediction (Top-1): komondor (sheep dog) (87.71%)

Actual Prediction (Top-1): **window shade (61.65%)**

Note: (1) Images are for illustration; actual images were cropped to 224x224 and mean-centered prior to model training and classification; (2) Probability/certainty can increase as well as decrease. For analysis purposes, absolute value error magnitude was considered.

Example: Probability Error
(Certainty Degradation)[(2)]

Correct Classification, Different Probability

Image: ILSVRC2012_val_00024059
Model: resnet-18 (SAT)

Golden Model Prediction (Top-1):
water ouzel, dipper (bird) (95.76%)

Actual Prediction:
water ouzel, dipper (74.71%)

Probability Error = **-21.06%**

Single events induced faults can impact prediction accuracy and certainty;
In addition to SEFI, SEE analysis should account for Datapath signatures

Versal™ Adaptive SoC AI/ML Proton Test Results

Mitigating Datapath Errors with Fault-Aware Training (FAT)

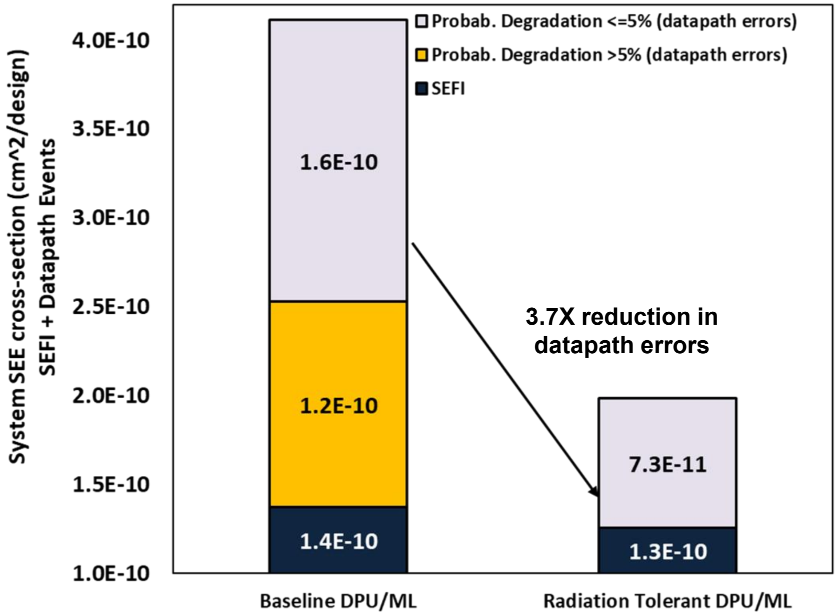

Versal AI Engine SEFI Cross-Section (per VC1902)

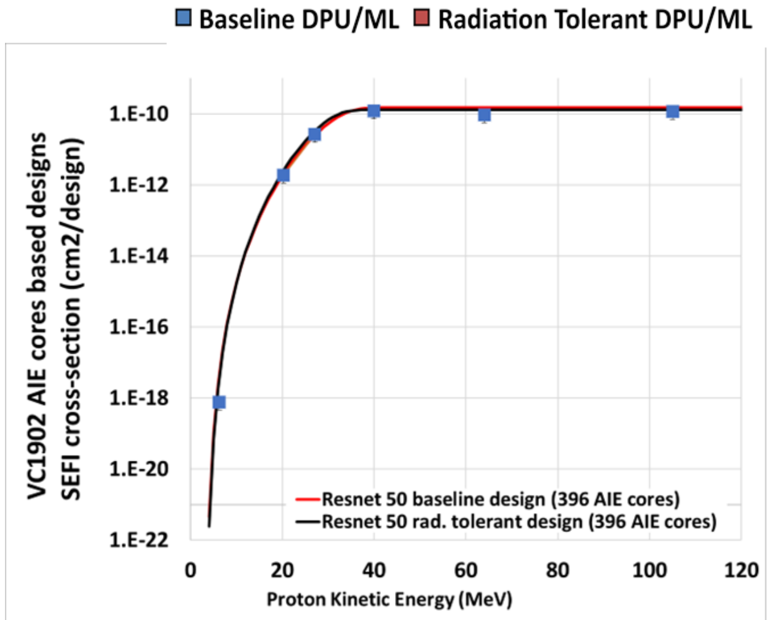

Radiation tolerant neural network response (vs. non-mitigated/baseline implementation), ResNet-50 network

- ~ 4X reduction in datapath SEU induced errors w/ rad. tolerant FAT platform (vs. baseline)
- > 5% probability degradation events are fully mitigated in the rad-tolerant design
- SEFI occurrence is <10% of the overall rad-tolerant platform single event cross-section
- 1.5 event per ~ 400 cores per year (*estimates using CREME96 AP8-Max; 500km, 52° inclination)*
- Weibull parameters published in Radiation Tolerant Versal AI Core Data Sheet (DS946)

P. Maillard et al., "Radiation-Tolerant Deep Learning Processor Unit (DPU)-Based Platform Using Xilinx 7-nm Versal Adaptive SoC AIE", IEEE NSREC 2024

AMD Versal™ Adaptive SoC Support Ecosystem

Versal™ Adaptive SoC Support Ecosystem

- Configuration Memory
  - 3D-Plus
  - Avalanche
  - DDC
  - Infineon
  - Mercury Systems
- Power Distribution
  - Frontgrade (CAES)
  - Infineon
  - Renesas
  - Texas Instruments
  - Vicor
- Development Platforms
  - Alpha Data ADK-VA601 (XCVC1902)
  - Alpha Data ADM-VB630 (XCVE2302)
  - iWave iW-RainboW-G57M® (XCVE2302)
  - Trenz Electronic TE0950-01-EGBE11A (XCVE2302)

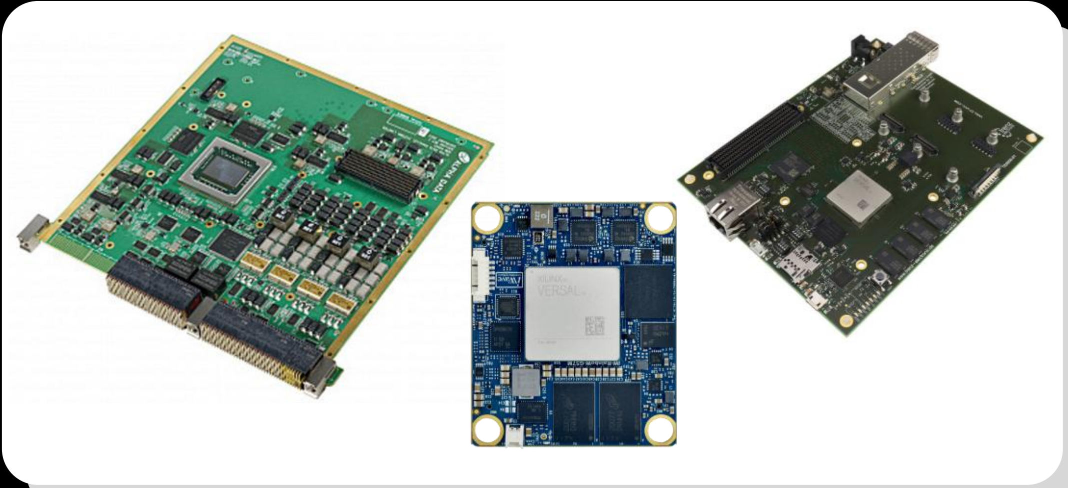

AMD Versal™ Adaptive SoC Design Examples and Conference Papers

Versal™ Adaptive SoCs – Design Examples, Conference Papers

RF Beamforming in AMD XQR Versal™ Adaptive SoCs using AI Engines

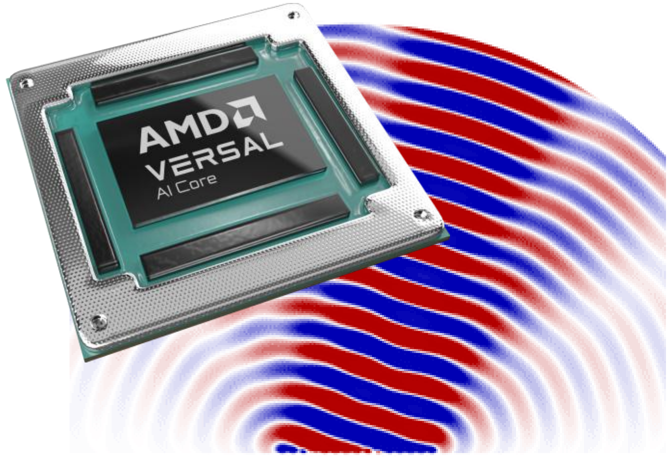

IEEE Space Computing Conference, July 2023
XQRVC1902

Radar Space Time Adaptive Processing using AMD Versal™ Adaptive SoCs

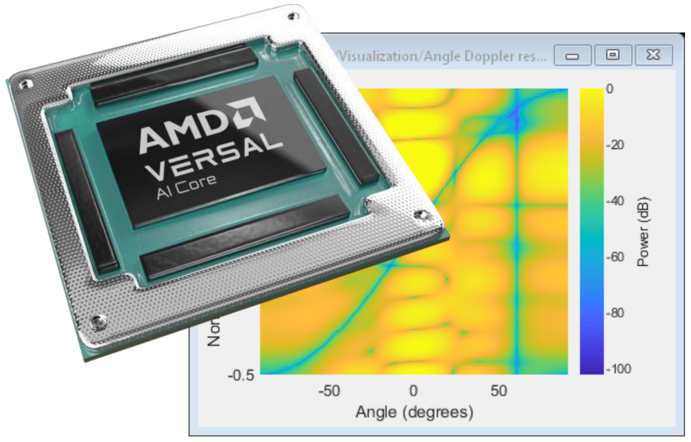

IEEE Space Computing Conference, July 2024
XQRVC1902

On-Orbit Anomaly Detection in Spacecraft Telemetry using RNNs in AMD Versal™ Adaptive SoCs

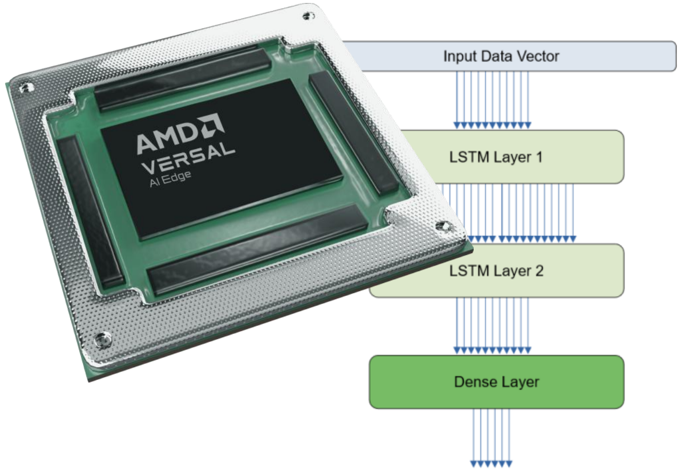

IEEE Space Computing Conference, July 2024
XQRVE2302